HP StoreAll 9300/9320 Storage Administrator Guide

Abstract

This guide describes tasks related to cluster configuration and monitoring, system upgrade and recovery, hardware component replacement, and troubleshooting for the HP 9300 Storage Gateway and the HP 9320 Storage. It does not document StoreAll file system features or standard Linux administrative tools and commands. For information about configuring and using StoreAll software file system features, see the *HP StoreAll Storage File System User Guide*.

This guide is intended for system administrators and technicians who are experienced with installing and administering networks, and with performing Linux operating and administrative tasks. For the latest StoreAll guides, browse to http://www.hp.com/support/StoreAllManuals.

HP Part Number: AW549-96068
Published: April 2013
Edition: 12

© Copyright 2010, 2013 Hewlett-Packard Development Company, L.P.
Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Acknowledgments

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Warranty

WARRANTY STATEMENT: To obtain a copy of the warranty for this product, see the warranty information website:

http://www.hp.com/go/storagewarranty

Revision History

Edition	Date	Software Version	Description
1	December 2009	5.3.1	Initial release of the 9300 Storage Gateway and 9320 Network Storage System administration guides.
2	April 2010	5.4	Added network management and support ticket.
3	August 2010	5.4.1	Added management console backup, migration to an agile management console configuration, software upgrade procedures, and system recovery procedures.
4	August 2010	5.4.1	Revised upgrade procedure.
5	December 2010	5.5	Added information about NDMP backups and configuring virtual interfaces, and updated cluster procedures.
6	March 2011	5.5	Updated segment evacuation information.
7	April 2011	5.6	Revised upgrade procedure and updated server information.
8	September 2011	6.0	Added or revised information about agile management console, NTP servers, Statistics tool, Ibrix Collect, event notification, upgrades.
9	June 2012	6.1	Combined the 9300 and 9320 administration guides, added or revised information about segment evacuation, events, Statistics tool, software upgrades, HP Insight Remote Support.
10	December 2012	6.2	Added or revised information about High Availability, failover, server tuning, segment migration and evacuation, SNMP, added upgrade checklist for common upgrade tasks.
11	March 2013	6.3	Updated information on upgrades, remote support, collection logs, phone home and troubleshooting. Now point users to website for the latest spare parts list instead of shipping the list. Added before and after upgrade steps for Express Query when going from 6.2 to 6.3.
12	April 2013	6.3	Removed post upgrade step that tells users to modify the `/etc/hosts` file on every StoreAll node. In the "Cascading Upgrades" appendix, added a section that tells users to ensure that the NFS exports option `subtree_check` is the default export option for every NFS export when upgrading from a StoreAll 5.x release. Also changed `ibrix_fm -m nofmfailover -A` to `ibrix_fm -m maintenance -A` in the "Cascading Upgrades" appendix. Updated information about SMB share creation.

Contents

1 Upgrading the StoreAll software to the 6.3 release....10
Online upgrades for StoreAll software....12
Preparing for the upgrade....12
Performing the upgrade....13
After the upgrade....13
Automated offline upgrades for StoreAll software 6.x to 6.3....14
Preparing for the upgrade....14
Performing the upgrade....14
After the upgrade....15
Manual offline upgrades for StoreAll software 6.x to 6.3....15
Preparing for the upgrade....15
Performing the upgrade manually....16
After the upgrade....17
Upgrading Linux StoreAll clients....18
Installing a minor kernel update on Linux clients....18
Upgrading Windows StoreAll clients....19
Upgrading pre-6.3 Express Query enabled file systems....19
Required steps before the StoreAll Upgrade....19
Required steps after the StoreAll Upgrade....20
Troubleshooting upgrade issues....21
Automatic upgrade....21
Manual upgrade....22
Offline upgrade fails because iLO firmware is out of date....22
Node is not registered with the cluster network....22
File system unmount issues....23
File system in MIF state after StoreAll software 6.3 upgrade....23
2 Product description....25
9300 Storage Gateway....25
9320 Storage System....25
System Components....25
HP StoreAll software features....25
High availability and redundancy....26
3 Getting started....27
Setting up the system....27
Installation steps....27
Additional configuration steps....27
Management interfaces....28
Using the StoreAll Management Console....29
Customizing the GUI....31
Adding user accounts for Management Console access....32
Using the CLI....32
Starting the array management software....32
StoreAll client interfaces....33
StoreAll software manpages....33
Changing passwords....33
Configuring ports for a firewall....34
Configuring NTP servers....35
Configuring HP Insight Remote Support on StoreAll systems....35
Configuring the StoreAll cluster for Insight Remote Support....37
Configuring Insight Remote Support for HP SIM 7.1 and IRS 5.7....39

Configuring Insight Remote Support for HP SIM 6.3 and IRS 5.6....42
Testing the Insight Remote Support configuration....45
Updating the Phone Home configuration....45
Disabling Phone Home....46
Troubleshooting Insight Remote Support....46
4 Configuring virtual interfaces for client access....48
Network and VIF guidelines....48
Creating a bonded VIF....49
Configuring backup servers....49
Configuring NIC failover....49
Configuring automated failover....50
Example configuration....50
Specifying VIFs in the client configuration....50
Configuring VLAN tagging....51
Configuring link state monitoring for iSCSI network interfaces....51
5 Configuring failover....53
Agile management consoles....53
Agile Fusion Manager modes....53
Viewing information about Fusion Managers....53
Agile Fusion Manager and failover....53
Configuring High Availability on the cluster....54
What happens during a failover....55
Configuring automated failover with the HA Wizard....55
Configuring automated failover manually....62
Changing the HA configuration manually....63
Failing a server over manually....64
Failing back a server....64
Setting up HBA monitoring....64
Checking the High Availability configuration....66
Capturing a core dump from a failed node....68
Prerequisites for setting up the crash capture....68
Setting up nodes for crash capture....69
6 Configuring cluster event notification....70
Cluster events....70
Setting up email notification of cluster events....70
Associating events and email addresses....71
Configuring email notification settings....71
Dissociating events and email addresses....71
Testing email addresses....71
Viewing email notification settings....72
Setting up SNMP notifications....72
Configuring the SNMP agent....72
Configuring trapsink settings....73
Associating events and trapsinks....74
Defining views....74
Configuring groups and users....74
Deleting elements of the SNMP configuration....75
Listing SNMP configuration information....75
Event notification for MSA array systems....75
7 Configuring system backups....77
Backing up the Fusion Manager configuration....77
Using NDMP backup applications....77
Configuring NDMP parameters on the cluster....78

NDMP process management....79
Viewing or canceling NDMP sessions....79
Starting, stopping, or restarting an NDMP Server....79
Viewing or rescanning tape and media changer devices....80
NDMP events....80
8 Creating host groups for StoreAll clients....81
How host groups work....81
Creating a host group tree....81
Adding a StoreAll client to a host group....82
Adding a domain rule to a host group....82
Viewing host groups....83
Deleting host groups....83
Other host group operations....83
9 Monitoring cluster operations....84
Monitoring 9300/9320 hardware....84
Monitoring servers....84
Monitoring hardware components....88
Obtaining server details....88
Monitoring storage and storage components....92
Managing LUNs in a storage cluster....93
Monitoring the status of file serving nodes....93
Monitoring cluster events....94
Viewing events....94
Removing events from the events database table....95
Monitoring cluster health....95
Health checks....96
Health check reports....96
Viewing logs....98
Viewing operating statistics for file serving nodes....98
10 Using the Statistics tool....100
Installing and configuring the Statistics tool....100
Installing the Statistics tool....100
Enabling collection and synchronization....100
Upgrading the Statistics tool from StoreAll software 6.0....101
Using the Historical Reports GUI....101
Generating reports....102
Deleting reports....103
Maintaining the Statistics tool....103
Space requirements....103
Updating the Statistics tool configuration....104
Changing the Statistics tool configuration....104
Fusion Manager failover and the Statistics tool configuration....104
Checking the status of Statistics tool processes....105
Controlling Statistics tool processes....105
Troubleshooting the Statistics tool....105
Log files....106
Uninstalling the Statistics tool....106
11 Maintaining the system....107
Shutting down the system....107
Shutting down the StoreAll software....107
Powering off the hardware....108
Starting the system....108
Starting the StoreAll software....108

Powering file serving nodes on or off....109
Performing a rolling reboot....109
Starting and stopping processes....110
Tuning file serving nodes and StoreAll clients....110
Managing segments....114
Migrating segments....115
Evacuating segments and removing storage from the cluster....118
Removing a node from a cluster....120
Maintaining networks....122
Cluster and user network interfaces....122
Adding user network interfaces....122
Setting network interface options in the configuration database....123
Preferring network interfaces....123
Unpreferring network interfaces....125
Making network changes....125
Changing the IP address for a Linux StoreAll client....125
Changing the IP address for the cluster interface on a dedicated management console....125
Changing the cluster interface....125
Managing routing table entries....126
Deleting a network interface....126
Viewing network interface information....126
12 Licensing....128
Viewing license terms....128
Retrieving a license key....128
Using AutoPass to retrieve and install permanent license keys....128
13 Upgrading firmware....129
Components for firmware upgrades....129
Steps for upgrading the firmware....130
Finding additional information on FMT....133
Downloading MSA2000 G2/G3 firmware for 9320 systems....133
14 Troubleshooting....134
Collecting information for HP Support with the IbrixCollect....134
Collecting logs....134
Downloading the archive file....135
Deleting the archive file....135
Configuring Ibrix Collect....136
Obtaining custom logging from ibrix_collect add-on scripts....137
Creating an add-on script....137
Running an add-on script....138
Viewing the output from an add-on script....138
Viewing data collection information....139
Adding/deleting commands or logs in the XML file....140
Viewing software version numbers....140
Troubleshooting specific issues....140
Software services....140
Failover....141
Windows StoreAll clients....141
Synchronizing information on file serving nodes and the configuration database....142
Troubleshooting an Express Query Manual Intervention Failure (MIF)....142
15 Recovering a file serving node....144
Obtaining the latest StoreAll software release....144
Performing the recovery....144
Completing the restore on a file serving node....149

The ibrix_auth command fails after a restore....150
16 Support and other resources....151
Contacting HP....151
Related information....151
Obtaining spare parts....152
HP websites....152
Rack stability....152
Product warranties....152
Subscription service....152
17 Documentation feedback....153
A Cascading Upgrades....154
Upgrading the StoreAll software to the 6.1 release....154
Online upgrades for StoreAll software 6.x to 6.1....155
Preparing for the upgrade....155
Performing the upgrade....155
After the upgrade....156
Offline upgrades for StoreAll software 5.6.x or 6.0.x to 6.1....156
Preparing for the upgrade....156
Performing the upgrade....157
After the upgrade....158
Upgrading Linux StoreAll clients....158
Installing a minor kernel update on Linux clients....159
Upgrading Windows StoreAll clients....159
Upgrading pre-6.0 file systems for software snapshots....159
Upgrading pre-6.1.1 file systems for data retention features....161
Troubleshooting upgrade issues....161
Automatic upgrade....161
Manual upgrade....162
Offline upgrade fails because iLO firmware is out of date....162
Node is not registered with the cluster network....162
File system unmount issues....163
Upgrading the StoreAll software to the 5.6 release....163
Automatic upgrades....164
Manual upgrades....164
Preparing for the upgrade....164
Saving the node configuration....165
Performing the upgrade....165
Restoring the node configuration....165
Completing the upgrade....165
Troubleshooting upgrade issues....166
Automatic upgrade....166
Manual upgrade....167
Upgrading the StoreAll software to the 5.5 release....167
Automatic upgrades....167
Manual upgrades....168
Standard upgrade for clusters with a dedicated Management Server machine or blade....168
Standard online upgrade....169
Standard offline upgrade....170
Agile upgrade for clusters with an agile management console configuration....172
Agile online upgrade....172
Agile offline upgrade....176
Troubleshooting upgrade issues....178

B Component diagrams for 9300 systems....180
Front view of file serving node....180
Rear view of file serving node....180
C System component and cabling diagrams for 9320 systems....183
System component diagrams....183
Front view of 9300c array controller or 9300cx 3.5" 12-drive enclosure....183
Rear view of 9300c array controller....184
Rear view of 9300cx 3.5" 12-drive enclosure....184
Front view of file serving node....185
Rear view of file serving node....185
Cabling diagrams....188
Cluster network cabling diagram....188
SATA option cabling....189
SAS option cabling....190
Drive enclosure cabling....191
D Warnings and precautions....192
Electrostatic discharge information....192
Preventing electrostatic discharge....192
Grounding methods....192
Equipment symbols....193
Rack warnings and precautions....193
Device warnings and precautions....194
E Regulatory compliance notices....196
Regulatory compliance identification numbers....196
Federal Communications Commission notice....196
FCC rating label....196
Class A equipment....196
Class B equipment....196
Modification....197
Cables....197
Canadian notice (Avis Canadien)....197
Class A equipment....197
Class B equipment....197
European Union notice....197
Japanese notices....198
Japanese VCCI-A notice....198
Japanese VCCI-B notice....198
Japanese VCCI marking....198
Japanese power cord statement....198
Korean notices....198
Class A equipment....198
Class B equipment....198
Taiwanese notices....199
BSMI Class A notice....199
Taiwan battery recycle statement....199
Turkish recycling notice....199
Vietnamese Information Technology and Communications compliance marking....199
Laser compliance notices....199
English laser notice....199
Dutch laser notice....200
French laser notice....200
German laser notice....200
Italian laser notice....201

Japanese laser notice....201
Spanish laser notice....201
Recycling notices....202
English recycling notice....202
Bulgarian recycling notice....202
Czech recycling notice....202
Danish recycling notice....202
Dutch recycling notice....202
Estonian recycling notice....203
Finnish recycling notice....203
French recycling notice....203
German recycling notice....203
Greek recycling notice....203
Hungarian recycling notice....203
Italian recycling notice....204
Latvian recycling notice....204
Lithuanian recycling notice....204
Polish recycling notice....204
Portuguese recycling notice....204
Romanian recycling notice....205
Slovak recycling notice....205
Spanish recycling notice....205
Swedish recycling notice....205
Battery replacement notices....206
Dutch battery notice....206
French battery notice....206
German battery notice....207
Italian battery notice....207
Japanese battery notice....208
Spanish battery notice....208
Glossary....209
Index....211

1 Upgrading the StoreAll software to the 6.3 release

This chapter describes how to upgrade to the 6.3 StoreAll software release.

IMPORTANT: Print the following table and check off each step as you complete it.

NOTE: (Upgrades from version 6.0.x) CIFS share permissions are granted on a global basis in v6.0.X. When upgrading from v6.0.X, confirm that the correct share permissions are in place.

Table 1 Prerequisites checklist for all upgrades

Step	Description	Step completed?
1	Verify that the entire cluster is currently running StoreAll 6.0 or later by entering the following command: `ibrix_version -l` **IMPORTANT:** All the StoreAll nodes must be at the same release. • If you are running a version of StoreAll earlier than 6.0, upgrade the product as described in "Cascading Upgrades" (page 154). • If you are running StoreAll 6.0 or later, proceed with the upgrade steps in this section.	☐
2	Verify that the `/local` partition contains at least 4 GB for the upgrade by using the following command: `df -kh /local`	☐
3	The 6.3 release requires that nodes hosting the agile Fusion Manager be registered on the cluster network. Run the following command to verify that nodes hosting the agile Fusion Manager have IP addresses on the cluster network: `ibrix_fm -l` If a node is configured on the user network, see "Node is not registered with the cluster network " (page 22) for a workaround. **NOTE:** The Fusion Manager and all file serving nodes must be upgraded to the new release at the same time. Do not change the active/passive Fusion Manager configuration during the upgrade.	☐

4	Verify that the crash kernel parameter on all nodes has been set to 256M by viewing the default boot entry in the `/etc/grub.conf` file, as shown in the following example:
	`kernel /vmlinuz-2.6.18-194.el5 ro root=/dev/vg1/lv1 crashkernel=256M@16M`
	The `/etc/grub.conf` file might contain multiple instances of the crash kernel parameter. Make sure you modify each instance that appears in the file.
	If you must modify the `/etc/grub.conf` file, follow the steps in this section:
	1. Use SSH to access the active Fusion Manager (FM).
	2. Do one of the following:
	• (Versions 6.2 and later) Place all passive FMs into nofmfailover mode:
	`ibrix_fm -m nofmfailover -A`
	• (Versions earlier than 6.2) Place all passive FMs into maintenance mode:
	`ibrix_fm -m maintenance -A`
	3. Disable Segment Server Failover on each node in the cluster:
	`ibrix_server -m -U -h <node>`
	4. Set the crash kernel to 256M in the `/etc/grub.conf` file. The `/etc/grub.conf` file might contain multiple instances of the crash kernel parameter. Make sure you modify each instance that appears in the file.
	NOTE: Save a copy of the `/etc/grub.conf` file before you modify it.
		☐

Table 1 Prerequisites checklist for all upgrades *(continued)*

Step	Description	Step completed?

	The following example shows the crash kernel set to 256M:
	`kernel /vmlinuz-2.6.18-194.el5 ro root=/dev/vg1/lv1 crashkernel=256M@16M`
	5. Reboot the active FM.
	6. Use SSH to access each passive FM and do the following:
	a. Modify the `/etc/grub.conf` file as described in the previous steps.
	b. Reboot the node.
	7. After all nodes in the cluster are back up, use SSH to access the active FM.
	8. Place all disabled FMs back into passive mode:
	`ibrix_fm -m passive -A`
	9. Re-enable Segment Server Failover on each node:
	`ibrix_server -m -h <node>`	

5	If your cluster includes G6 servers, check the iLO2 firmware version. This issue does not affect G7 servers. The firmware must be at version 2.05 for HA to function properly. If your servers have an earlier version of the iLO2 firmware, run the `CP014256.scexe` script as described in the following steps:
	1. Make sure the `/local/ibrix` folder is empty prior to copying the contents of pkgfull. When you upgrade the StoreAll software later in this chapter, this folder must contain only `.rpm` packages listed in the build manifest for the upgrade or the upgrade will fail.
	2. Mount the ISO image and copy the entire directory structure to the `/local/ibrix` directory.
	The following is an example of the mount command:
	`mount -o loop /local/pkg/ibrix-pkgfull-FS_6.3.72+IAS_6.3.72-x86_64.signed.iso /mnt/<storeall>`
	In this example, *<storeall>* can have any name.
	The following is an example of the copy command:
	`cp -R /mnt/storeall/* /local/ibrix`
	3. Execute the firmware binary at the following location:
	`/local/ibrix/distrib/firmware/CP014256.scexe`
		☐

6	Make sure StoreAll is running the latest firmware. For information on how to find the version of firmware that StoreAll is running, see the Administrator Guide for your release.	☐
7	If you are using 1GBe with mode 6, consider switching to mode 4. See the HP StoreAll Storage Best Practices Guide for additional information.	☐

8	Verify that all file system nodes can "see" and "access" every segment logical volume that the file system node is configured for as either the owner or the backup by entering the following commands:
	1. To view all segments, logical volume name, and owner, enter the following command on one line:
	`ibrix_fs -i | egrep -e OWNER -e MIXED|awk '{ print $1, $3, $6, $2, $14, $5}' | tr " " "\t"`
	2. To verify the visibility of the correct segments on the current file system node enter the following command on each file system node:
	`lvm lvs | awk '{print $1}'`
		☐

9	Ensure that no active tasks are running. Stop any active remote replication, data tiering, or rebalancer tasks running on the cluster. (Use `ibrix_task -l` to list active tasks.) When the upgrade is complete, you can start the tasks again.	☐

Table 1 Prerequisites checklist for all upgrades *(continued)*

Step	Description	Step completed?
	For additional information on how to stop a task, enter the `ibrix_task` command for the help.	
10	Record all host tunings, FS tunings and FS mounting options by using the following commands: 1. To display file system tunings, enter: `ibrix_fs_tune -l >/local/ibrix_fs_tune-l.txt` 2. To display default StoreAll tunings and settings, enter: `ibrix_host_tune -L >/local/ibrix_host_tune-L.txt` 3. To display all non-default configuration tunings and settings, enter: `ibrix_host_tune -q >/local/ibrix_host_tune-q.txt`	☐

11	Ensure that the "ibrix" local user account exists and it has the same UID number on all the servers in the cluster. If they do not have the same UID number, create the account and change the UIDs as needed to make them the same on all the servers. Similarly, ensure that the "ibrix-user" local user group exists and has the same GID number on all servers.
	Enter the following commands on each node:
	`grep ibrix /etc/passwd`
	`grep ibrix-user /etc/group`
		☐

12	Ensure that all nodes are up and running. To determine the status of your cluster nodes, check the health of each server by either using the dashboard on the Management Console or entering the `ibrix_health -i -h <hostname>` command for each node in the cluster. At the top of the output look for "`PASSED`."	☐
13	If you have one or more Express Query enabled file system, each one needs to be manually upgraded as described in "Upgrading pre-6.3 Express Query enabled file systems" (page 19). **IMPORTANT:** Run the steps in "Required steps before the StoreAll Upgrade" (page 19) before the upgrade. This section provides steps for saving your custom metadata and audit log. After you upgrade the StoreAll software, run the steps in "Required steps after the StoreAll Upgrade" (page 20). These post-upgrade steps are required for you to preserve your custom metadata and audit log data.	☐

Online upgrades for StoreAll software

Online upgrades are supported only from the StoreAll 6.x release. Upgrades from earlier StoreAll releases must use the appropriate offline upgrade procedure.

When performing an online upgrade, note the following:

- File systems remain mounted and client I/O continues during the upgrade.
- The upgrade process takes approximately 45 minutes, regardless of the number of nodes.
- The total I/O interruption per node IP is four minutes, allowing for a failover time of two minutes and a failback time of two additional minutes.
- Client I/O having a timeout of more than two minutes is supported.

Preparing for the upgrade

To prepare for the upgrade, complete the following steps, ensure that high availability is enabled on each node in the cluster by running the following command:

```
ibrix_haconfig -l
```

If the command displays an `Overall HA Configuration Checker Results - PASSED` status, high availability is enabled on each node in the cluster. If the command returns `Overall`

`HA Configuration Checker Results - FAILED`, complete the following list items based on the result returned for each component:

1. Make sure you have completed all steps in the upgrade checklist (Table 1 (page 10)).
2. If `Failed` was displayed for the HA Configuration or Auto Failover columns or both, perform the steps described in the section "Configuring High Availability on the cluster" in the administrator guide for your current release.
3. If `Failed` was displayed for the NIC or HBA Monitored columns, see the sections for `ibrix_nic -m -h <host> -A node_2/node_interface` and `ibrix_hba -m -h <host> -p <World_Wide_Name>` in the CLI guide for your current release.

Performing the upgrade

The online upgrade is supported only from the StoreAll 6.x releases.

IMPORTANT: Complete all steps provided in the Table 1 (page 10).

Complete the following steps:

1. StoreAll OS version 6.3 is only available through the registered release process. To obtain the ISO image, contact HP Support to register for the release and obtain access to the software dropbox.
2. Make sure the `/local/ibrix` folder is empty prior to copying the contents of `pkgfull`. The upgrade will fail if the `/local/ibrix` folder contains leftover `.rpm` packages not listed in the build manifest.
3. Mount the ISO image and copy the entire directory structure to the `/local/ibrix` directory on the disk running the OS.

   The following is an example of the mount command:

   ```
   mount -o loop
   /local/pkg/ibrix-pkgfull-FS_6.3.72+IAS_6.3.72-x86_64.signed.iso
   /mnt/<storeall>
   ```

   In this example, *<storeall>* can have any name.

   The following is an example of the copy command:

   ```
   cp -R /mnt/storeall/* /local/ibrix
   ```
4. Change directory to `/local/ibrix` and then run `chmod -R 777 *` on the entire directory structure.
5. Run the upgrade script and follow the on-screen directions:

   ```
   ./auto_online_ibrixupgrade
   ```
6. Upgrade Linux StoreAll clients. See "Upgrading Linux StoreAll clients" (page 18).
7. If you received a new license from HP, install it as described in "Licensing" (page 128).

After the upgrade

Complete these steps:

1. If your cluster nodes contain any 10Gb NICs, reboot these nodes to load the new driver. You must do this step before you upgrade the server firmware, as requested later in this procedure.
2. Upgrade your firmware as described in "Upgrading firmware" (page 129).
3. Start any remote replication, rebalancer, or data tiering tasks that were stopped before the upgrade.
4. If you have a file system version prior to version 6, you might have to make changes for snapshots and data retention, as mentioned in the following list:
   - **Snapshots.** Files used for snapshots must either be created on StoreAll software 6.0 or later, or the pre-6.0 file system containing the files must be upgraded for snapshots. To

upgrade a file system, use the `upgrade60.sh` utility. For more information, see "Upgrading pre-6.0 file systems for software snapshots" (page 159).

- **Data retention.** Files used for data retention (including WORM and auto-commit) must be created on StoreAll software 6.1.1 or later, or the pre-6.1.1 file system containing the files must be upgraded for retention features. To upgrade a file system, use the `ibrix_reten_adm -u -f FSNAME` command. Additional steps are required before and after you run the `ibrix_reten_adm -u -f FSNAME` command. For more information, see "Upgrading pre-6.0 file systems for software snapshots" (page 159).

5. If you have an Express Query enabled file system prior to version 6.3, manually complete each file system upgrade as described in "Required steps after the StoreAll Upgrade" (page 20).

Automated offline upgrades for StoreAll software 6.x to 6.3

Preparing for the upgrade

To prepare for the upgrade, complete the following steps:

1. Make sure you have completed all steps in the upgrade checklist (Table 1 (page 10)).
2. Stop all client I/O to the cluster or file systems. On the Linux client, use `lsof </mountpoint>` to show open files belonging to active processes.
3. Verify that all StoreAll file systems can be successfully unmounted from all FSN servers:

   ```
   ibrix_umount -f fsname
   ```

Performing the upgrade

This upgrade method is supported only for upgrades from StoreAll software 6.x to the 6.3 release. Complete the following steps:

1. StoreAll OS version 6.3 is only available through the registered release process. To obtain the ISO image, contact HP Support to register for the release and obtain access to the software dropbox.
2. Make sure the `/local/ibrix` folder is empty prior to copying the contents of `pkgfull`. The upgrade will fail if the `/local/ibrix` folder contains leftover `.rpm` packages not listed in the build manifest.
3. Mount the ISO image and copy the entire directory structure to the `/local/ibrix` directory on the disk running the OS.

   The following is an example of the mount command:

   ```
   mount -o loop
   /local/pkg/ibrix-pkgfull-FS_6.3.72+IAS_6.3.72-x86_64.signed.iso
   /mnt/<storeall>
   ```

   In this example, `<storeall>` can have any name.

   The following is an example of the copy command:

   ```
   cp -R /mnt/storeall/* /local/ibrix
   ```

4. Change directory to `/local/ibrix` on the disk running the OS and then run `chmod -R 777 *` on the entire directory structure.
5. Run the following upgrade script:

   ```
   ./auto_ibrixupgrade
   ```

   The upgrade script automatically stops the necessary services and restarts them when the upgrade is complete. The upgrade script installs the Fusion Manager on all file serving nodes. The Fusion Manager is in active mode on the node where the upgrade was run, and is in passive mode on the other file serving nodes. If the cluster includes a dedicated Management Server, the Fusion Manager is installed in passive mode on that server.

6. Upgrade Linux StoreAll clients. See "Upgrading Linux StoreAll clients" (page 18).
7. If you received a new license from HP, install it as described in the "Licensing" chapter in this guide.

After the upgrade

Complete the following steps:

1. If your cluster nodes contain any 10Gb NICs, reboot these nodes to load the new driver. You must do this step before you upgrade the server firmware, as requested later in this procedure.
2. Upgrade your firmware as described in "Upgrading firmware" (page 129).
3. Mount file systems on Linux StoreAll clients.
4. If you have a file system version prior to version 6, you might have to make changes for snapshots and data retention, as mentioned in the following list:
   - **Snapshots.** Files used for snapshots must either be created on StoreAll software 6.0 or later, or the pre-6.0 file system containing the files must be upgraded for snapshots. To upgrade a file system, use the `upgrade60.sh` utility. For more information, see "Upgrading pre-6.0 file systems for software snapshots" (page 159).
   - **Data retention.** Files used for data retention (including WORM and auto-commit) must be created on StoreAll software 6.1.1 or later, or the pre-6.1.1 file system containing the files must be upgraded for retention features. To upgrade a file system, use the `ibrix_reten_adm -u -f` *`FSNAME`* command. Additional steps are required before and after you run the `ibrix_reten_adm -u -f` *`FSNAME`* command. For more information, see "Upgrading pre-6.0 file systems for software snapshots" (page 159).
5. If you have an Express Query enabled file system prior to version 6.3, manually complete each file system upgrade as described in "Required steps after the StoreAll Upgrade" (page 20).

Manual offline upgrades for StoreAll software 6.x to 6.3

Preparing for the upgrade

To prepare for the upgrade, complete the following steps:

1. Make sure you have completed all steps in the upgrade checklist (Table 1 (page 10)).
2. Verify that `ssh` shared keys have been set up. To do this, run the following command on the node hosting the active instance of the agile Fusion Manager:

   `ssh <server_name>`

   Repeat this command for each node in the cluster.
3. Verify that all file system node servers have separate file systems mounted on the following partitions by using the `df` command:
   - `/`
   - `/local`
   - `/stage`
   - `/alt`
4. Verify that all FSN servers have a minimum of 4 GB of free/available storage on the `/local` partition by using the `df` command .
5. Verify that all FSN servers are not reporting any partition as 100% full (at least 5% free space) by using the `df` command .
6. Note any custom tuning parameters, such as file system mount options. When the upgrade is complete, you can reapply the parameters.
7. Stop all client I/O to the cluster or file systems. On the Linux client, use `lsof </mountpoint>` to show open files belonging to active processes.

8. On the active Fusion Manager, enter the following command to place the Fusion Manager into maintenance mode:

   ```
   <ibrixhome>/bin/ibrix_fm -m nofmfailover -P -A
   ```

9. On the active Fusion Manager node, disable automated failover on all file serving nodes:

   ```
   <ibrixhome>/bin/ibrix_server -m -U
   ```

10. Run the following command to verify that automated failover is off. In the output, the HA column should display `off`.

    ```
    <ibrixhome>/bin/ibrix_server -l
    ```

11. Unmount file systems on Linux StoreAll clients:

    ```
    ibrix_umount -f MOUNTPOINT
    ```

12. Stop the SMB, NFS and NDMP services on all nodes. Run the following commands on the node hosting the active Fusion Manager:

    ```
    ibrix_server -s -t cifs -c stop
    ibrix_server -s -t nfs -c stop
    ibrix_server -s -t ndmp -c stop
    ```

    If you are using SMB, verify that all likewise services are down on all file serving nodes:

    ```
    ps -ef | grep likewise
    ```

    Use `kill -9` to stop any likewise services that are still running.

    If you are using NFS, verify that all NFS processes are stopped:

    ```
    ps -ef | grep nfs
    ```

    If necessary, use the following command to stop NFS services:

    ```
    /etc/init.d/nfs stop
    ```

    Use `kill -9` to stop any NFS processes that are still running.

    If necessary, run the following command on all nodes to find any open file handles for the mounted file systems:

    ```
    lsof </mountpoint>
    ```

    Use `kill -9` to stop any processes that still have open file handles on the file systems.

13. Unmount each file system manually:

    ```
    ibrix_umount -f FSNAME
    ```

    Wait up to 15 minutes for the file systems to unmount.

    Troubleshoot any issues with unmounting file systems before proceeding with the upgrade. See "File system unmount issues" (page 23).

Performing the upgrade manually

This upgrade method is supported only for upgrades from StoreAll software 6.x to the 6.3 release. Complete the following steps:

1. StoreAll OS version 6.3 is only available through the registered release process. To obtain the ISO image, contact HP Support to register for the release and obtain access to the software dropbox.
2. Make sure the `/local/ibrix` folder is empty prior to copying the contents of `pkgfull`. The upgrade will fail if the `/local/ibrix` folder contains leftover `.rpm` packages not listed in the build manifest.
3. Mount the ISO image on each node and copy the entire directory structure to the `/local/ibrix` directory on the disk running the OS.

   The following is an example of the mount command:

```
mount -o loop
/local/pkg/ibrix-pkgfull-FS_6.3.72+IAS_6.3.72-x86_64.signed.iso
/mnt/<storeall>
```

In this example, *<storeall>* can have any name.

The following is an example of the copy command:

```
cp -R /mnt/storeall/* /local/ibrix
```

4. Change directory to `/local/ibrix` on the disk running the OS and then run `chmod -R 777 *` on the entire directory structure.
5. Run the following upgrade script:

   ```
   ./ibrixupgrade –f
   ```

   The upgrade script automatically stops the necessary services and restarts them when the upgrade is complete. The upgrade script installs the Fusion Manager on all file serving nodes. The Fusion Manager is in active mode on the node where the upgrade was run, and is in passive mode on the other file serving nodes. If the cluster includes a dedicated Management Server, the Fusion Manager is installed in passive mode on that server.
6. Upgrade Linux StoreAll clients. See "Upgrading Linux StoreAll clients" (page 18).
7. If you received a new license from HP, install it as described in the "Licensing" chapter in this guide.

After the upgrade

Complete the following steps:

1. If your cluster nodes contain any 10Gb NICs, reboot these nodes to load the new driver. You must do this step before you upgrade the server firmware, as requested later in this procedure.
2. Upgrade your firmware as described in "Upgrading firmware" (page 129).
3. Run the following command to rediscover physical volumes:

   ```
   ibrix_pv -a
   ```
4. Apply any custom tuning parameters, such as mount options.
5. Remount all file systems:

   ```
   ibrix_mount -f <fsname> -m </mountpoint>
   ```
6. Re-enable High Availability if used:

   ```
   ibrix_server -m
   ```
7. Start any remote replication, rebalancer, or data tiering tasks that were stopped before the upgrade.
8. If you are using SMB, set the following parameters to synchronize the SMB software and the Fusion Manager database:

   - `smb signing enabled`
   - `smb signing required`
   - `ignore_writethru`

   Use `ibrix_cifsconfig` to set the parameters, specifying the value appropriate for your cluster (1=enabled, 0=disabled). The following examples set the parameters to the default values for the 6.3 release:

   ```
   ibrix_cifsconfig -t -S "smb_signing_enabled=0,
   smb_signing_required=0"
   ```

   ```
   ibrix_cifsconfig -t -S "ignore_writethru=1"
   ```

   The SMB signing feature specifies whether clients must support SMB signing to access SMB shares. See the *HP StoreAll Storage File System User Guide* for more information about this

feature. When `ignore_writethru` is enabled, StoreAll software ignores writethru buffering to improve SMB write performance on some user applications that request it.

9. Mount file systems on Linux StoreAll clients.
10. If you have a file system version prior to version 6, you might have to make changes for snapshots and data retention, as mentioned in the following list:
    - **Snapshots.** Files used for snapshots must either be created on StoreAll software 6.0 or later, or the pre-6.0 file system containing the files must be upgraded for snapshots. To upgrade a file system, use the `upgrade60.sh` utility. For more information, see "Upgrading pre-6.0 file systems for software snapshots" (page 159).
    - **Data retention.** Files used for data retention (including WORM and auto-commit) must be created on StoreAll software 6.1.1 or later, or the pre-6.1.1 file system containing the files must be upgraded for retention features. To upgrade a file system, use the `ibrix_reten_adm -u -f` *`FSNAME`* command. Additional steps are required before and after you run the `ibrix_reten_adm -u -f` *`FSNAME`* command. For more information, see "Upgrading pre-6.0 file systems for software snapshots" (page 159).
11. If you have an Express Query enabled file system prior to version 6.3, manually complete each file system upgrade as described in "Required steps after the StoreAll Upgrade" (page 20).

Upgrading Linux StoreAll clients

Be sure to upgrade the cluster nodes before upgrading Linux StoreAll clients. Complete the following steps on each client:

1. Download the latest HP StoreAll client 6.3 package.
2. Expand the tar file.
3. Run the upgrade script:

   ```
   ./ibrixupgrade -tc -f
   ```

   The upgrade software automatically stops the necessary services and restarts them when the upgrade is complete.
4. Execute the following command to verify the client is running StoreAll software:

   ```
   /etc/init.d/ibrix_client status
   IBRIX Filesystem Drivers loaded
   IBRIX IAD Server (pid 3208) running...
   ```

The `IAD` service should be running, as shown in the previous sample output. If it is not, contact HP Support.

Installing a minor kernel update on Linux clients

The StoreAll client software is upgraded automatically when you install a compatible Linux minor kernel update.

If you are planning to install a minor kernel update, first run the following command to verify that the update is compatible with the StoreAll client software:

`/usr/local/ibrix/bin/verify_client_update` *`<kernel_update_version>`*

The following example is for a RHEL 4.8 client with kernel version `2.6.9-89.ELsmp`:

```
# /usr/local/ibrix/bin/verify_client_update 2.6.9-89.35.1.ELsmp
Kernel update 2.6.9-89.35.1.ELsmp is compatible.
```

If the minor kernel update is compatible, install the update with the vendor RPM and reboot the system. The StoreAll client software is then automatically updated with the new kernel, and StoreAll client services start automatically. Use the `ibrix_version -l -C` command to verify the kernel version on the client.

NOTE: To use the `verify_client` command, the StoreAll client software must be installed.

Upgrading Windows StoreAll clients

Complete the following steps on each client:

1. Remove the old Windows StoreAll client software using the **Add or Remove Programs** utility in the Control Panel.
2. Copy the Windows StoreAll client MSI file for the upgrade to the machine.
3. Launch the Windows Installer and follow the instructions to complete the upgrade.
4. Register the Windows StoreAll client again with the cluster and check the option to **Start Service after Registration**.
5. Check **Administrative Tools** | **Services** to verify that the StoreAll client service is started.
6. Launch the Windows StoreAll client. On the Active Directory Settings tab, click **Update** to retrieve the current Active Directory settings.
7. Mount file systems using the StoreAll Windows client GUI.

NOTE: If you are using Remote Desktop to perform an upgrade, you must log out and log back in to see the drive mounted.

Upgrading pre-6.3 Express Query enabled file systems

The internal database schema format of Express Query enabled file systems changed between releases 6.2.x and 6.3. Each file system with Express Query enabled must be manually upgraded to 6.3. This section has instructions to be run before and after the StoreAll upgrade, on each of those file systems.

Required steps before the StoreAll Upgrade

These steps are required before the StoreAll Upgrade:

1. Mount all Express Query file systems on the cluster to be upgraded if they are not mounted yet.
2. Save your custom metadata by entering the following command:

   ```
   /usr/local/ibrix/bin/MDExport.pl --dbconfig
   /usr/local/Metabox/scripts/startup.xml --database <FSNAME>
   --outputfile /tmp/custAttributes.csv --user ibrix
   ```

3. Save your audit log data by entering the following commands:

   ```
   ibrix_audit_reports -t time -f <FSNAME>
   ```

   ```
   cp <path to report file printed from previous command>
   /tmp/auditData.csv
   ```

4. Disable auditing by entering the following command:

   ```
   ibrix_fs -A -f <FSNAME> -oa audit_mode=off
   ```

   In this instance *<FSNAME>* is the file system.

5. If any archive API shares exist for the file system, delete them.

NOTE: To list all HTTP shares, enter the following command:

```
ibrix_httpshare -l
```

To list only REST API (Object API) shares, enter the following command:

```
ibrix_httpshare -l -f <FSNAME> -v 1 | grep "objectapi: true" | awk
'{ print $2 }'
```

In this instance *<FSNAME>* is the file system.

- Delete all HTTP shares, regular or REST API (Object API) by entering the following command:

  ```
  ibrix_httpshare -d -f <FSNAME>
  ```

  In this instance *<FSNAME>* is the file system.

- Delete a specific REST API (Object API) share by entering the following command:

  ```
  ibrix_httpshare -d <SHARENAME> -c <PROFILENAME> -t <VHOSTNAME>
  ```

  In this instance

  - *<SHARENAME>* is the share name.
  - *<PROFILENAME>* is the profile name.
  - *<VHOSTNAME>* is the virtual host name

6. Disable Express Query by entering the following command:

   ```
   ibrix_fs -T -D -f <FSNAME>
   ```

7. Shut down Archiving daemons for Express Query by entering the following command:

   ```
   ibrix_archiving -S -F
   ```

8. delete the internal database files for this file system by entering the following command:

   ```
   rm -rf <FS_MOUNTPOINT>/.archiving/database
   ```

   In this instance *<FS_MOUNTPOINT>* is the file system mount point.

Required steps after the StoreAll Upgrade

These steps are required after the StoreAll Upgrade:

1. Restart the Archiving daemons for Express Query:
2. Re-enable Express Query on the file systems you disabled it from before by entering the following command:

   ```
   ibrix_fs -T -E -f <FSNAME>
   ```

   In this instance *<FSNAME>* is the file system.

   Express Query will begin resynchronizing (repopulating) a new database for this filesystem.

3. Re-enable auditing if you had it running before (the default) by entering the following command:

   ```
   ibrix_fs -A -f <FSNAME> -oa audit_mode=on
   ```

   In this instance *<FSNAME>* is the file system.

4. Re-create REST API (Object API) shares deleted before the upgrade on each node in the cluster (if desired) by entering the following command:

NOTE: The REST API (Object API) functionality has expanded, and any REST API (Object API) shares you created in previous releases are now referred to as HTTP-StoreAll REST API shares in file-compatible mode. The 6.3 release is also introducing a new type of share called HTTP-StoreAll REST API share in Object mode.

```
ibrix_httpshare -a <SHARENAME> -c <PROFILENAME> -t <VHOSTNAME> -f
<FSNAME> -p <DIRPATH> -P <URLPATH> -S
“ibrixRestApiMode=filecompatible, anonymous=true”
```

In this instance:

- *<SHARENAME>* is the share name.
- *<PROFILENAME>* is the profile name.
- *<VHOSTNAME>* is the virtual host name
- *<FSNAME>* is the file system.
- *<DIRPATH>* is the directory path.
- *<URLPATH>* is the URL path.
- *<SETTINGLIST>* is the settings.

5. Wait for the resynchronizer to complete by entering the following command until its output is `<FSNAME>: OK`:

   ```
   ibrix_archiving -l
   ```

6. Restore your audit log data by entering the following command:

   ```
   MDImport -f <FSNAME> -n /tmp/auditData.csv -t audit
   ```

   In this instance *<FSNAME>* is the file system.

7. Restore your custom metadata by entering the following command:

   ```
   MDImport -f <FSNAME> -n /tmp/custAttributes.csv -t custom
   ```

   In this instance *<FSNAME>* is the file system.

Troubleshooting upgrade issues

If the upgrade does not complete successfully, check the following items. For additional assistance, contact HP Support.

Automatic upgrade

Check the following:

- If the initial execution of `/usr/local/ibrix/setup/upgrade` fails, check `/usr/local/ibrix/setup/upgrade.log` for errors. It is imperative that all servers are up and running the StoreAll software before you execute the upgrade script.
- If the install of the new OS fails, power cycle the node. Try rebooting. If the install does not begin after the reboot, power cycle the machine and select the `upgrade` line from the `grub` boot menu.
- After the upgrade, check `/usr/local/ibrix/setup/logs/postupgrade.log` for errors or warnings.
- If configuration restore fails on any node, look at `/usr/local/ibrix/autocfg/logs/appliance.log` on that node to determine which feature restore failed. Look at the specific feature log file under `/usr/local/ibrix/setup/logs/` for more detailed information.

  To retry the copy of configuration, use the following command:

  ```
  /usr/local/ibrix/autocfg/bin/ibrixapp upgrade -f -s
  ```

- If the install of the new image succeeds, but the configuration restore fails and you need to revert the server to the previous install, run the following command and then reboot the machine. This step causes the server to boot from the old version (the alternate partition).

  ```
  /usr/local/ibrix/setup/boot_info -r
  ```

- If the public network interface is down and inaccessible for any node, power cycle that node.

NOTE: Each node stores its `ibrixupgrade.log` file in `/tmp`.

Manual upgrade

Check the following:

- If the restore script fails, check `/usr/local/ibrix/setup/logs/restore.log` for details.
- If configuration restore fails, look at `/usr/local/ibrix/autocfg/logs/appliance.log` to determine which feature restore failed. Look at the specific feature log file under `/usr/local/ibrix/setup/logs/` for more detailed information.

  To retry the copy of configuration, use the following command:

  ```
  /usr/local/ibrix/autocfg/bin/ibrixapp upgrade -f -s
  ```

Offline upgrade fails because iLO firmware is out of date

If the iLO2 firmware is out of date on a node, the `auto_ibrixupgrade` script will fail. The `/usr/local/ibrix/setup/logs/auto_ibrixupgrade.log` reports the failure and describes how to update the firmware.

After updating the firmware, run the following command on the node to complete the StoreAll software upgrade:

```
/local/ibrix/ibrixupgrade -f
```

Node is not registered with the cluster network

Nodes hosting the agile Fusion Manager must be registered with the cluster network. If the `ibrix_fm` command reports that the IP address for a node is on the user network, you will need to reassign the IP address to the cluster network. For example, the following commands report that node `ib51-101`, which is hosting the active Fusion Manager, has an IP address on the user network (`192.168.51.101`) instead of the cluster network.

```
[root@ib51-101 ibrix]# ibrix_fm -i
FusionServer: ib51-101 (active, quorum is running)
================================================
[root@ib51-101 ibrix]# ibrix_fm -l
NAME      IP ADDRESS
--------  ----------
ib51-101  192.168.51.101
ib51-102  10.10.51.102
```

1. If the node is hosting the active Fusion Manager, as in this example, stop the Fusion Manager on that node:

   ```
   [root@ib51-101 ibrix]# /etc/init.d/ibrix_fusionmanager stop
   Stopping Fusion Manager Daemon                          [  OK  ]
   [root@ib51-101 ibrix]#
   ```

2. On the node now hosting the active Fusion Manager (`ib51-102` in the example), unregister node `ib51-101`:

```
[root@ib51-102 ~]# ibrix_fm -u ib51-101
Command succeeded!
```

3. On the node hosting the active Fusion Manager, register node `ib51-101` and assign the correct IP address:

```
[root@ib51-102 ~]# ibrix_fm -R ib51-101 -I 10.10.51.101
Command succeeded!
```

NOTE: When registering a Fusion Manager, be sure the hostname specified with `-R` matches the hostname of the server.

The `ibrix_fm` commands now show that node `ib51-101` has the correct IP address and node `ib51-102` is hosting the active Fusion Manager.

```
[root@ib51-102 ~]# ibrix_fm -f
NAME      IP ADDRESS
--------  ----------
ib51-101  10.10.51.101
ib51-102  10.10.51.102
[root@ib51-102 ~]# ibrix_fm -i
FusionServer: ib51-102 (active, quorum is running)
================================================
```

File system unmount issues

If a file system does not unmount successfully, perform the following steps on all servers:

1. Run the following commands:

```
chkconfig ibrix_server off
chkconfig ibrix_ndmp off
chkconfig ibrix_fusionmanager off
```

2. Reboot all servers.
3. Run the following commands to move the services back to the `on` state. The commands do not start the services.

```
chkconfig ibrix_server on
chkconfig ibrix_ndmp on
chkconfig ibrix_fusionmanager on
```

4. Run the following commands to start the services:

```
service ibrix_fusionmanager start
service ibrix_server start
```

5. Unmount the file systems and continue with the upgrade procedure.

File system in MIF state after StoreAll software 6.3 upgrade

If an Express Query enabled file systems ended in MIF state after completing the StoreAll software upgrade process (`ibrix_archiving -l prints <FSNAME>: MIF`), check the MIF status by running the following command:

```
cat /<FSNAME>/.archiving/database/serialization/ManualInterventionFailure
```

If the command's output displays `Version mismatch, upgrade needed` (as shown in the following output), steps were not performed as described in "Required steps after the StoreAll Upgrade" (page 20).

```
MIF:Version mismatch, upgrade needed. (error code 14)
```

If you did not see the `Version mismatch, upgrade needed` in the command's output, see "Troubleshooting an Express Query Manual Intervention Failure (MIF)" (page 142).

Perform the following steps only if you see the `Version mismatch, upgrade needed` in the command's output:

1. Disable auditing by entering the following command:
   ```
   ibrix_fs -A -f <FSNAME> -oa audit_mode=off
   ```
   In this instance *<FSNAME>* is the file system.
2. Disable Express Query by entering the following command:
   ```
   ibrix_fs -T -D -f <FSNAME>
   ```
   In this instance *<FSNAME>* is the file system.
3. Delete the internal database files for this file system by entering the following command:
   ```
   rm -rf <FS_MOUNTPOINT>/.archiving/database
   ```
   In this instance *<FS_MOUNTPOINT>* is the file system mount point.
4. Clear the MIF condition by running the following command:
   ```
   ibrix_archiving -C <FSNAME>
   ```
   In this instance *<FSNAME>* is the file system.
5. Re-enable Express Query on the file systems:
   ```
   ibrix_fs -T -E -f <FSNAME>
   ```
   In this instance *<FSNAME>* is the file system.
   Express Query will begin resynchronizing (repopulating) a new database for this file system.
6. Re-enable auditing if you had it running before (the default).
   ```
   ibrix_fs -A -f <FSNAME> -oa audit_mode=on
   ```
   In this instance *<FSNAME>* is the file system.
7. Restore your audit log data:
   ```
   MDImport -f <FSNAME> -n /tmp/auditData.csv -t audit
   ```
   In this instance *<FSNAME>* is the file system.
8. Restore your custom metadata:
   ```
   MDImport -f <FSNAME> -n /tmp/custAttributes.csv -t custom
   ```
   In this instance *<FSNAME>* is the file system.

2 Product description

This guide provides information about configuring, monitoring, and maintaining HP StoreAll 9300 Storage Gateways and 9320 Storage.

IMPORTANT: It is important to keep regular backups of the cluster configuration.

9300 Storage Gateway

The 9300 Storage Gateway is a flexible, scale-out solution that brings gateway file services to HP MSA, EVA, P4000, or 3rd-party arrays or SANs. The system provides the following features:

- Segmented, scalable file system under a single namespace
- NFS, SMB (Server Message Block), FTP, and HTTP support for accessing file system data
- Centralized CLI and GUI cluster management
- Policy management
- Continuous remote replication

9320 Storage System

The 9320 Storage System is a highly available, scale-out storage solution for file data workloads. The system combines HP StoreAll File Serving Software with HP server and storage hardware to create an expansible cluster of file serving nodes. The system provides the following features:

- Segmented, scalable file system under a single namespace
- NFS, SMB, FTP, and HTTP support for accessing file system data
- Centralized CLI and GUI cluster management
- Policy management
- Continuous remote replication
- Dual redundant paths to all storage components
- Gigabytes-per-second throughput

System Components

For 9300 system components, see "Component diagrams for 9300 systems" (page 180).

For 9320 system components, see "System component and cabling diagrams for 9320 systems" (page 183).

For a complete list of system components, see the HP StoreAll Storage System QuickSpecs, which are available at:

http://www.hp.com/go/StoreAll

HP StoreAll software features

HP StoreAll software is a scale-out, network-attached storage solution including a parallel file system for clusters, an integrated volume manager, high-availability features such as automatic failover of multiple components, and a centralized management interface. StoreAll software can scale to thousands of nodes.

Based on a segmented file system architecture, StoreAll software integrates I/O and storage systems into a single clustered environment that can be shared across multiple applications and managed from a central Fusion Manager.

StoreAll software is designed to operate with high-performance computing applications that require high I/O bandwidth, high IOPS throughput, and scalable configurations.

Some of the key features and benefits are as follows:

- Scalable configuration. You can add servers to scale performance and add storage devices to scale capacity.
- Single namespace. All directories and files are contained in the same namespace.
- Multiple environments. Operates in both the SAN and DAS environments.
- High availability. The high-availability software protects servers.
- Tuning capability. The system can be tuned for large or small-block I/O.
- Flexible configuration. Segments can be migrated dynamically for rebalancing and data tiering.

High availability and redundancy

The segmented architecture is the basis for fault resilience—loss of access to one or more segments does not render the entire file system inaccessible. Individual segments can be taken offline temporarily for maintenance operations and then returned to the file system.

To ensure continuous data access, StoreAll software provides manual and automated failover protection at various points:

- **Server.** A failed node is powered down and a designated standby server assumes all of its segment management duties.
- **Segment.** Ownership of each segment on a failed node is transferred to a designated standby server.
- **Network interface.** The IP address of a failed network interface is transferred to a standby network interface until the original network interface is operational again.
- **Storage connection.** For servers with HBA-protected Fibre Channel access, failure of the HBA triggers failover of the node to a designated standby server.

3 Getting started

IMPORTANT: Follow these guidelines when using your system:

- Do not modify any parameters of the operating system or kernel, or update any part of the 9320 Storage unless instructed to do so by HP; otherwise, the system could fail to operate properly.
- File serving nodes are tuned for file serving operations. With the exception of supported backup programs, do not run other applications directly on the nodes.

Setting up the system

Installation steps

An HP service specialist sets up the system at your site, including the following tasks:

- Remove the product from the shipping cartons that you have placed in the location where the product will be installed, confirm the contents of each carton against the list of included items and check for any physical damage to the exterior of the product, and connect the product to the power and network provided by you.
- Review your server, network, and storage environment relevant to the HP Enterprise NAS product implementation to validate that prerequisites have been met.
- Validate that your file system performance, availability, and manageability requirements have not changed since the service planning phase. Finalize the HP Enterprise NAS product implementation plan and software configuration.
- Implement the documented and agreed-upon configuration based on the information you provided on the pre-delivery checklist.
- Document configuration details.

Additional configuration steps

When your system is up and running, you can continue configuring the cluster and file systems. The Management Console GUI and CLI are used to perform most operations. (Some features described here may be configured for you as part of the system installation.)

Cluster. Configure the following as needed:

- Firewall ports. See "Configuring ports for a firewall" (page 34)
- HP Insight Remote Support and Phone Home. See "Configuring HP Insight Remote Support on StoreAll systems" (page 35).
- Virtual interfaces for client access. See "Configuring virtual interfaces for client access" (page 48).
- Cluster event notification through email or SNMP. See "Configuring cluster event notification" (page 70).
- Fusion Manager backups. See "Backing up the Fusion Manager configuration" (page 77).
- NDMP backups. See "Using NDMP backup applications" (page 77).
- Statistics tool. See "Using the Statistics tool" (page 100).
- Ibrix Collect. See "Collecting information for HP Support with the IbrixCollect" (page 134).

File systems. Set up the following features as needed:

- NFS, SMB (Server Message Block), FTP, or HTTP. Configure the methods you will use to access file system data.
- Quotas. Configure user, group, and directory tree quotas as needed.
- Remote replication. Use this feature to replicate changes in a source file system on one cluster to a target file system on either the same cluster or a second cluster.
- Data retention and validation. Use this feature to manage WORM and retained files.
- Antivirus support. This feature is used with supported Antivirus software, allowing you to scan files on a StoreAll file system.
- StoreAll software snapshots. This feature allows you to capture a point-in-time copy of a file system or directory for online backup purposes and to simplify recovery of files from accidental deletion. Users can access the file system or directory as it appeared at the instant of the snapshot.
- Block Snapshots. This feature uses the array capabilities to capture a point-in-time copy of a file system for online backup purposes and to simplify recovery of files from accidental deletion. The snapshot replicates all file system entities at the time of capture and is managed exactly like any other file system.
- File allocation. Use this feature to specify the manner in which segments are selected for storing new files and directories.
- Data tiering. Use this feature to move files to specific tiers based on file attributes.

For more information about these file system features, see the *HP StoreAll Storage File System User Guide*.

Localization support

Red Hat Enterprise Linux 5 uses the UTF-8 (8-bit Unicode Transformation Format) encoding for supported locales. This allows you to create, edit and view documents written in different locales using UTF-8. StoreAll software supports modifying the `/etc/sysconfig/i18n` configuration file for your locale. The following example sets the LANG and SUPPORTED variables for multiple character sets:

```
LANG="ko_KR.utf8"
  SUPPORTED="en_US.utf8:en_US:en:ko_KR.utf8:ko_KR:ko:zh_CN.utf8:zh_CN:zh"
  SYSFONT="lat0-sun16"
  SYSFONTACM="iso15"
```

Management interfaces

Cluster operations are managed through the StoreAll Fusion Manager, which provides both a Management Console and a CLI. Most operations can be performed from either the StoreAll Management Console or the CLI.

The following operations can be performed only from the CLI:

- SNMP configuration (`ibrix_snmpagent`, `ibrix_snmpgroup`, `ibrix_snmptrap`, `ibrix_snmpuser`, `ibrix_snmpview`)
- Health checks (`ibrix_haconfig`, `ibrix_health`, `ibrix_healthconfig`)
- Raw storage management (`ibrix_pv`, `ibrix_vg`, `ibrix_lv`)
- Fusion Manager operations (`ibrix_fm`) and Fusion Manager tuning (`ibrix_fm_tune`)
- File system checks (`ibrix_fsck`)
- Kernel profiling (`ibrix_profile`)

- Cluster configuration (`ibrix_clusterconfig`)
- Configuration database consistency (`ibrix_dbck`)
- Shell task management (`ibrix_shell`)

The following operations can be performed only from the StoreAll Management Console:

- Scheduling recurring data validation scans
- Scheduling recurring software snapshots
- Scheduling recurring block snapshots

Using the StoreAll Management Console

The StoreAll Management Console is a browser-based interface to the Fusion Manager. See the release notes for the supported browsers and other software required to view charts on the dashboard. You can open multiple Management Console windows as necessary.

If you are using HTTP to access the Management Console, open a web browser and navigate to the following location, specifying port 80:

`http://<management_console_IP>:80/fusion`

If you are using HTTPS to access the Management Console, navigate to the following location, specifying port 443:

`https://<management_console_IP>:443/fusion`

In these URLs, *<management_console_IP>* is the IP address of the Fusion Manager user VIF.

The Management Console prompts for your user name and password. The default administrative user is **`ibrix`**. Enter the password that was assigned to this user when the system was installed. (You can change the password using the Linux `passwd` command.) To allow other users to access the Management Console, see "Adding user accounts for Management Console access" (page 32).

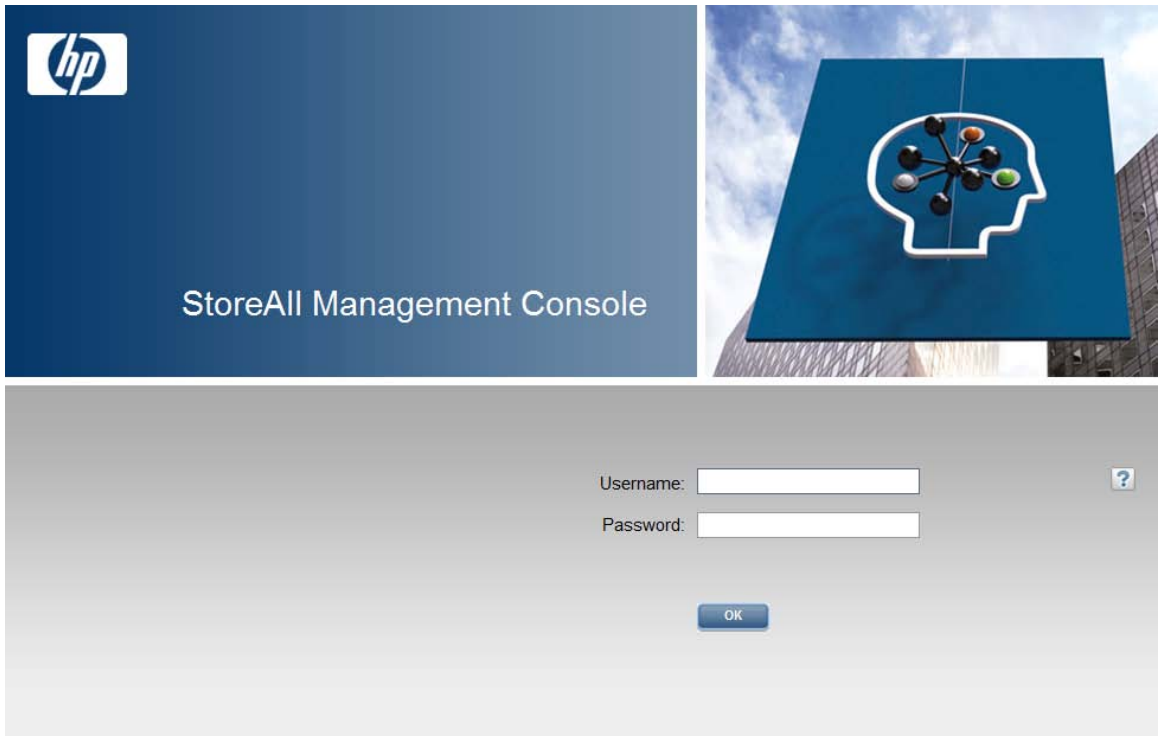

Upon login, the Management Console dashboard opens, allowing you to monitor the entire cluster. (See the online help for information about all Management Console displays and operations.) There are three parts to the dashboard: System Status, Cluster Overview, and the Navigator.

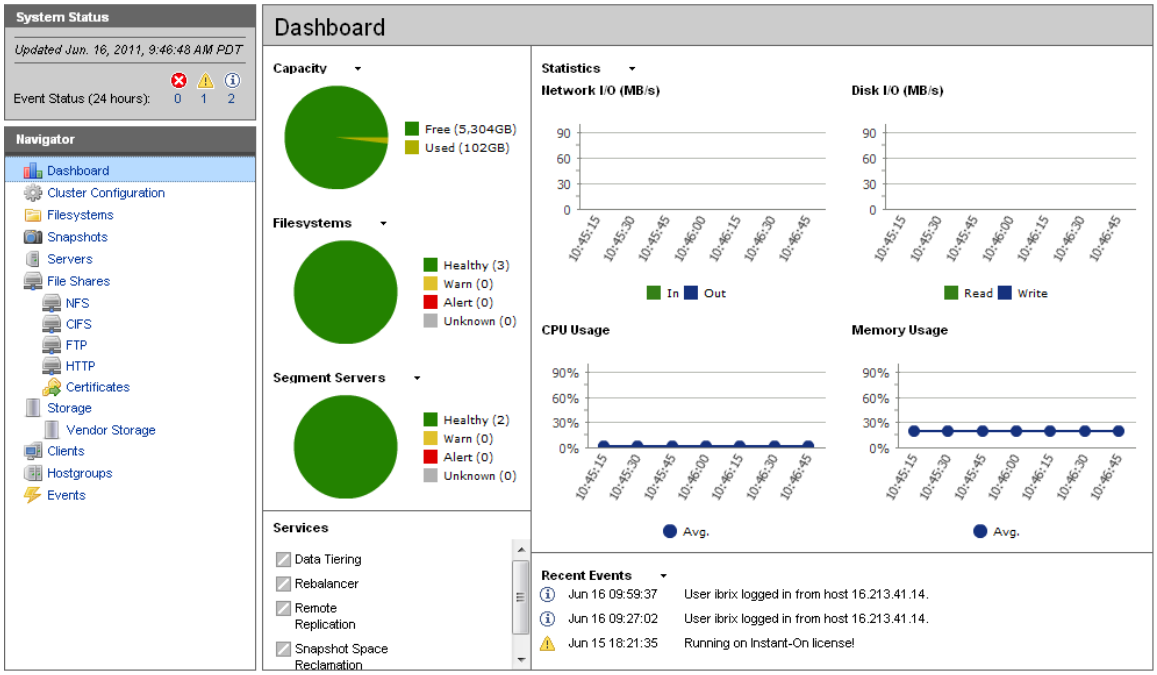

System Status

The System Status section lists the number of cluster events that have occurred in the last 24 hours. There are three types of events:

	Alerts. Disruptive events that can result in loss of access to file system data. Examples are a segment that is unavailable or a server that cannot be accessed.
	Warnings. Potentially disruptive conditions where file system access is not lost, but if the situation is not addressed, it can escalate to an alert condition. Examples are a very high server CPU utilization level or a quota limit close to the maximum.
	Information. Normal events that change the cluster. Examples are mounting a file system or creating a segment.

Cluster Overview

The Cluster Overview provides the following information:

Capacity

The amount of cluster storage space that is currently free or in use.

File systems

The current health status of the file systems in the cluster. The overview reports the number of file systems in each state (healthy, experiencing a warning, experiencing an alert, or unknown).

Segment Servers

The current health status of the file serving nodes in the cluster. The overview reports the number of nodes in each state (healthy, experiencing a warning, experiencing an alert, or unknown).

Services

Whether the specified file system services are currently running:

	One or more tasks are running.
	No tasks are running.

Statistics

Historical performance graphs for the following items:

- Network I/O (MB/s)
- Disk I/O (MB/s)
- CPU usage (%)
- Memory usage (%)

On each graph, the X-axis represents time and the Y-axis represents performance.

Use the **Statistics** menu to select the servers to monitor (up to two), to change the maximum value for the Y-axis, and to show or hide resource usage distribution for CPU and memory.

Recent Events

The most recent cluster events. Use the **Recent Events** menu to select the type of events to display.

You can also access certain menu items directly from the Cluster Overview. Mouse over the Capacity, Filesystems or Segment Server indicators to see the available options.

Navigator

The Navigator appears on the left side of the window and displays the cluster hierarchy. You can use the Navigator to drill down in the cluster configuration to add, view, or change cluster objects such as file systems or storage, and to initiate or view tasks such as snapshots or replication. When you select an object, a details page shows a summary for that object. The lower Navigator allows you to view details for the selected object, or to initiate a task. In the following example, we selected Filesystems in the upper Navigator and Mountpoints in the lower Navigator to see details about the mounts for file system `ifs1`.

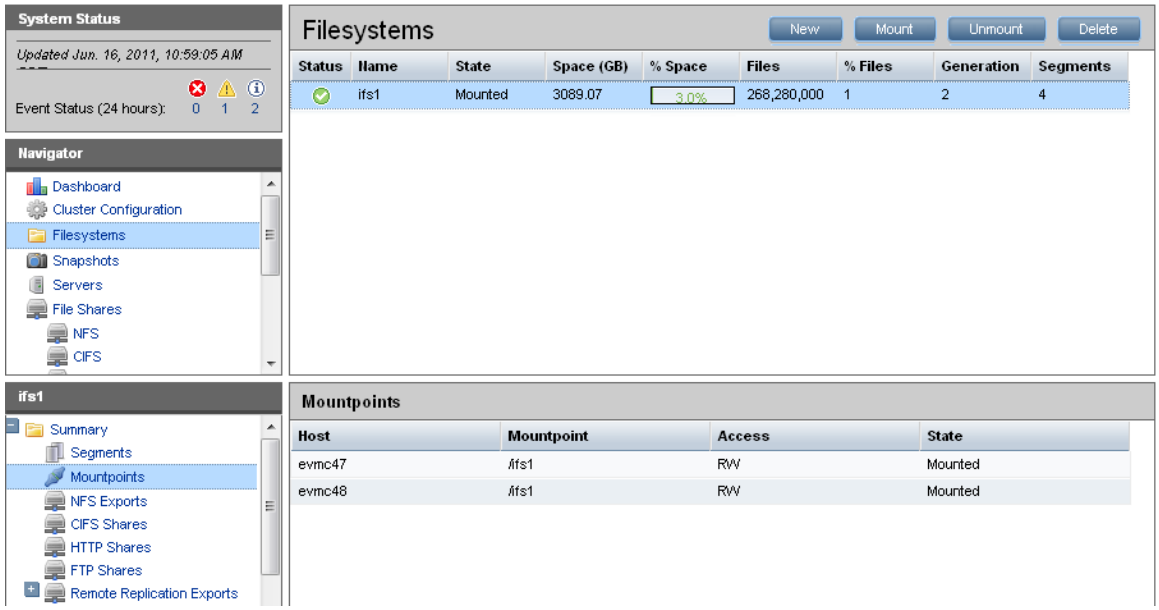

NOTE: When you perform an operation on the GUI, a spinning finger is displayed until the operation is complete. However, if you use Windows Remote Desktop to access the GUI, the spinning finger is not displayed.

Customizing the GUI

For most tables in the GUI, you can specify the columns that you want to display and the sort order of each column. When this feature is available, mousing over a column causes the label to change color and a pointer to appear. Click the pointer to see the available options. In the following

example, you can sort the contents of the Mountpoint column in ascending or descending order, and you can select the columns that you want to appear in the display.

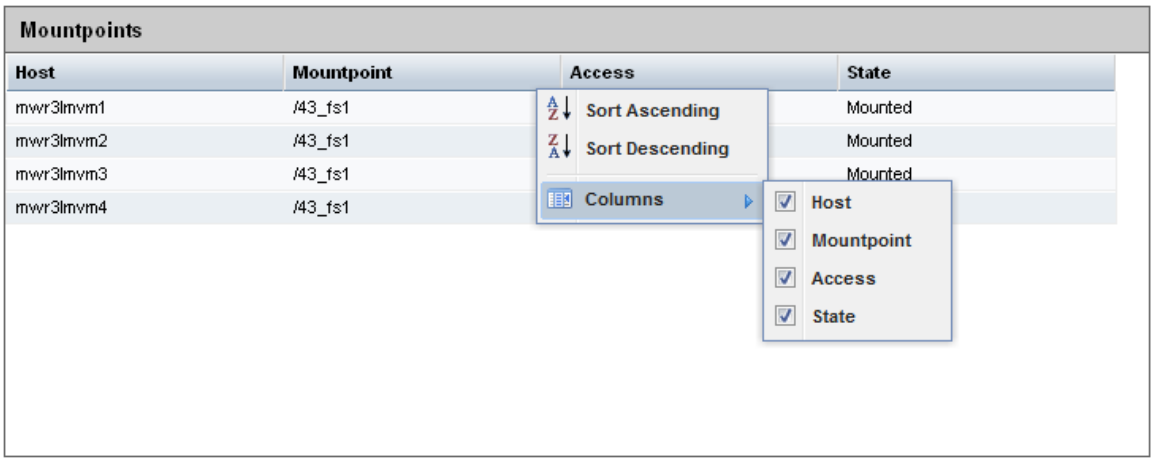

Adding user accounts for Management Console access

StoreAll software supports administrative and user roles. When users log in under the administrative role, they can configure the cluster and initiate operations such as remote replication or snapshots. When users log in under the user role, they can view the cluster configuration and status, but cannot make configuration changes or initiate operations. The default administrative user name is `ibrix`. The default regular username is `ibrixuser`.

User names for the administrative and user roles are defined in the `/etc/group` file. Administrative users are specified in the `ibrix-admin` group, and regular users are specified in the `ibrix-user` group. These groups are created when StoreAll software is installed. The following entries in the `/etc/group` file show the default users in these groups:

```
ibrix-admin:x:501:root,ibrix
ibrix-user:x:502:ibrix,ibrixUser,ibrixuser
```

You can add other users to these groups as needed, using Linux procedures. For example:

```
adduser -G ibrix-<groupname> <username>
```

When using the `adduser` command, be sure to include the `-G` option.

Using the CLI

The administrative commands described in this guide must be executed on the Fusion Manager host and require root privileges. The commands are located in `$IBRIXHOME/bin`. For complete information about the commands, see the *HP StoreAll Network Storage System CLI Reference Guide*.

When using `ssh` to access the machine hosting the Fusion Manager, specify the IP address of the Fusion Manager user VIF.

Starting the array management software

Depending on the array type, you can launch the array management software from the GUI. In the Navigator, select **Vendor Storage**, select your array from the Vendor Storage page, and click **Launch Storage Management**.

StoreAll client interfaces

StoreAll clients can access the Fusion Manager as follows:

- **Linux clients.** Use Linux client commands for tasks such as mounting or unmounting file systems and displaying statistics. See the *HP StoreAll Storage CLI Reference Guide* for details about these commands.
- **Windows clients.** Use the Windows client GUI for tasks such as mounting or unmounting file systems and registering Windows clients.

Using the Windows StoreAll client GUI

The Windows StoreAll client GUI is the client interface to the Fusion Manager. To open the GUI, double-click the desktop icon or select the StoreAll client program from the Start menu on the client. The client program contains tabs organized by function.

NOTE: The Windows StoreAll client GUI can be started only by users with Administrative privileges.

- **Status.** Shows the client's Fusion Manager registration status and mounted file systems, and provides access to the IAD log for troubleshooting.
- **Registration.** Registers the client with the Fusion Manager, as described in the *HP StoreAll Storage Installation Guide*.
- **Mount.** Mounts a file system. Select the Cluster Name from the list (the cluster name is the Fusion Manager name), enter the name of the file system to mount, select a drive, and then click **Mount**. (If you are using Remote Desktop to access the client and the drive letter does not appear, log out and log in again.)
- **Umount.** Unmounts a file system.
- **Tune Host.** Tunable parameters include the NIC to prefer (the client uses the cluster interface by default unless a different network interface is preferred for it), the communications protocol (UDP or TCP), and the number of server threads to use.
- **Active Directory Settings.** Displays current Active Directory settings.

For more information, see the client GUI online help.

StoreAll software manpages

StoreAll software provides manpages for most of its commands. To view the manpages, set the `MANPATH` variable to include the path to the manpages and then export it. The manpages are in the `$IBRIXHOME/man` directory. For example, if `$IBRIXHOME` is `/usr/local/ibrix` (the default), set the `MANPATH` variable as follows and then export the variable:

```
MANPATH=$MANPATH:/usr/local/ibrix/man
```

Changing passwords

You can change the following passwords on your system:

- **Hardware passwords.** See the documentation for the specific hardware for more information.
- **Root password.** Use the `passwd(8)` command on each server.
- **StoreAll software user password.** This password is created during installation and is used to log in to the GUI. The default is `ibrix`. You can change the password using the Linux `passwd` command.

  ```
  # passwd ibrix
  ```

  You will be prompted to enter the new password.

Configuring ports for a firewall

IMPORTANT: To avoid unintended consequences, HP recommends that you configure the firewall during scheduled maintenance times.

When configuring a firewall, you should be aware of the following:

- SELinux should be disabled.
- By default, NFS uses random port numbers for operations such as mounting and locking. These ports must be fixed so that they can be listed as exceptions in a firewall configuration file. For example, you will need to lock specific ports for `rpc.statd`, `rpc.lockd`, `rpc.mountd`, and `rpc.quotad`.
- It is best to allow all ICMP types on all networks; however, you can limit ICMP to types 0, 3, 8, and 11 if necessary.

Be sure to open the ports listed in the following table.

Port	Description
22/tcp	SSH
123/tcp, 123/upd	NTP
5353/udp	Multicast DNS, 224.0.0.251
12865/tcp	`netperf` tool
80/tcp 443/tcp	Fusion Manager to file serving nodes
5432/tcp 8008/tcp 9002/tcp 9005/tcp 9008/tcp 9009/tcp 9200/tcp	Fusion Manager and StoreAll file system
	Between file serving nodes and NFS clients (user network)
2049/tcp, 2049/udp	NFS
111/tcp, 111/udp	RPC
875/tcp, 875/udp	quota
32803/tcp	lockmanager
32769/udp	lockmanager
892/tcp, 892/udp	mount daemon
662/tcp, 662/udp	stat
2020/tcp, 2020/udp	stat outgoing
4000:4003/tcp	reserved for use by a custom application (CMU) and can be disabled if not used
137/udp 138/udp 139/tcp 445/tcp	Between file serving nodes and SMB clients (user network)
9000:9002/tcp 9000:9200/udp	Between file serving nodes and StoreAll clients (user network)

Port	Description
20/tcp, 20/udp 21/tcp, 21/udp	Between file serving nodes and FTP clients (user network)
7777/tcp 8080/tcp	Between GUI and clients that need to access the GUI
5555/tcp, 5555/udp	Dataprotector
631/tcp, 631/udp	Internet Printing Protocol (IPP)
1344/tcp, 1344/udp	ICAP

Configuring NTP servers

When the cluster is initially set up, primary and secondary NTP servers are configured to provide time synchronization with an external time source. The list of NTP servers is stored in the Fusion Manager configuration. The active Fusion Manager node synchronizes its time with the external source. The other file serving nodes synchronize their time with the active Fusion Manager node. In the absence of an external time source, the local hardware clock on the agile Fusion Manager node is used as the time source. This configuration method ensures that the time is synchronized on all cluster nodes, even in the absence of an external time source.

On StoreAll clients, the time is not synchronized with the cluster nodes. You will need to configure NTP servers on StoreAll clients.

List the currently configured NTP servers:

```
ibrix_clusterconfig -i -N
```

Specify a new list of NTP servers:

```
ibrix_clusterconfig -c -N SERVER1[,...,SERVERn]
```

Configuring HP Insight Remote Support on StoreAll systems

IMPORTANT: In the StoreAll software 6.1 release, the default port for the StoreAll SNMP agent changed from 5061 to 161. This port number cannot be changed.

NOTE: Configuring Phone Home enables the hp-snmp-agents service internally. As a result, a large number of error messages, such as the following, could occasionally appear in `/var/log/hp-snmp-agents/cma.log`:

```
Feb 08 13:05:54 x946s1 cmahostd[25579]: cmahostd: Can't update OS filesys
object: /ifs1 (PEER3023)
```

The cmahostd daemon is part of the hp-snmp-agents service. This error message occurs because the file system exceeds <n> TB. If this occurs, HP recommends that before you perform operations such as unmounting a file system or stopping services on a file serving node (using the `ibrix_server` command), you disable the hp-snmp-agent service on each server first:

```
service hp-snmp-agents stop
```

After remounting the file system or restarting services on the file serving node, restart the hp-snmp-agents service on each server:

```
service hp-snmp-agents start
```

Prerequisites

The required components for supporting StoreAll systems are preinstalled on the file serving nodes. You must install HP Insight Remote Support on a separate Windows system termed the Central Management Server (CMS):

- HP Insight Manager (HP SIM). This software manages HP systems and is the easiest and least expensive way to maximize system uptime and health.
- Insight Remote Support Advanced (IRSA). This version is integrated with HP Systems Insight Manager (SIM). It provides comprehensive remote monitoring, notification/advisories, dispatch, and proactive service support. IRSA and HP SIM together are referred to as the CMS.
- The Phone Home configuration does not support backup or standby NICs that are used for NIC failover. If backup NICs are currently configured, remove the backup NICs from all nodes before configuring Phone Home. After a successful Phone Home configuration, you can reconfigure the backup NICs.

The following versions of the software are supported.

- HP SIM 6.3 and IRSA 5.6
- HP SIM 7.1 and IRSA 5.7

IMPORTANT: Keep in mind the following:

- For each file serving node, add the physical user network interfaces (by entering the `ibrix_nic` command or selecting the **Server** > **NICs** tab in the GUI) so the interfaces can communicate with HP SIM.
- Ensure that all user network interfaces on each file serving node can communicate with the CMS.

IMPORTANT: Insight Remote Support Standard (IRSS) is not supported with StoreAll software 6.1 and later.

For product descriptions and information about downloading the software, see the HP Insight Remote Support Software web page:

http://www.hp.com/go/insightremotesupport

For information about HP SIM:

http://www.hp.com/products/systeminsightmanager

For IRSA documentation:

http://www.hp.com/go/insightremoteadvanced-docs

IMPORTANT: You must compile and manually register the StoreAll MIB file by using HP Systems Insight Manager:

1. Download `ibrixMib.txt` from `/usr/local/ibrix/doc/`.
2. Rename the file to `ibrixMib.mib`.
3. In HP Systems Insight Manager, complete the following steps:
   a. Unregister the existing MIB by entering the following command:

      ```
      <BASE>\mibs>mxmib -d ibrixMib.mib
      ```

   b. Copy the `ibrixMib.mib` file to the `<BASE>\mibs` directory, and then enter the following commands:

      ```
      <BASE>\mibs>mcompile ibrixMib.mib
      <BASE>\mibs>mxmib -a ibrixMib.cfg
      ```

For more information about the MIB, see the "Compiling and customizing MIBs" chapter in the *HP Systems Insight Manager User Guide*, which is available at:

http://www.hp.com/go/insightmanagement/sim/

Click **Support & Documents** and then click **Manuals**. Navigate to the user guide.

Limitations

Note the following:

- For StoreAll systems, the HP Insight Remote Support implementation is limited to hardware events.

Configuring the StoreAll cluster for Insight Remote Support

To enable 9300/9320 systems for remote support, first register MSA disk arrays and then configure Phone Home settings. All nodes in the cluster should be up when you perform this step.

NOTE: Configuring Phone Home removes any previous StoreAll `snmp` configuration details and populates the SNMP configuration with Phone Home configuration details. When Phone Home is enabled, you cannot use `ibrix_snmpagent` to edit or change the `snmp` agent configuration. However, you can use `ibrix_snmptrap` to add trapsink IPs and you can use `ibrix_event` to associate events to the trapsink IPs.

Registering MSA disk arrays

To register an MSA disk array with the cluster, run the following command:

```
# ibrix_vs -r -n STORAGENAME -t msa -I IP(s) -U USERNAME [-P PASSWORD]
```

Configuring Phone Home settings

To configure Phone Home on the GUI, select **Cluster Configuration** in the upper Navigator and then select **Phone Home** in the lower Navigator. The Phone Home Setup panel shows the current configuration.

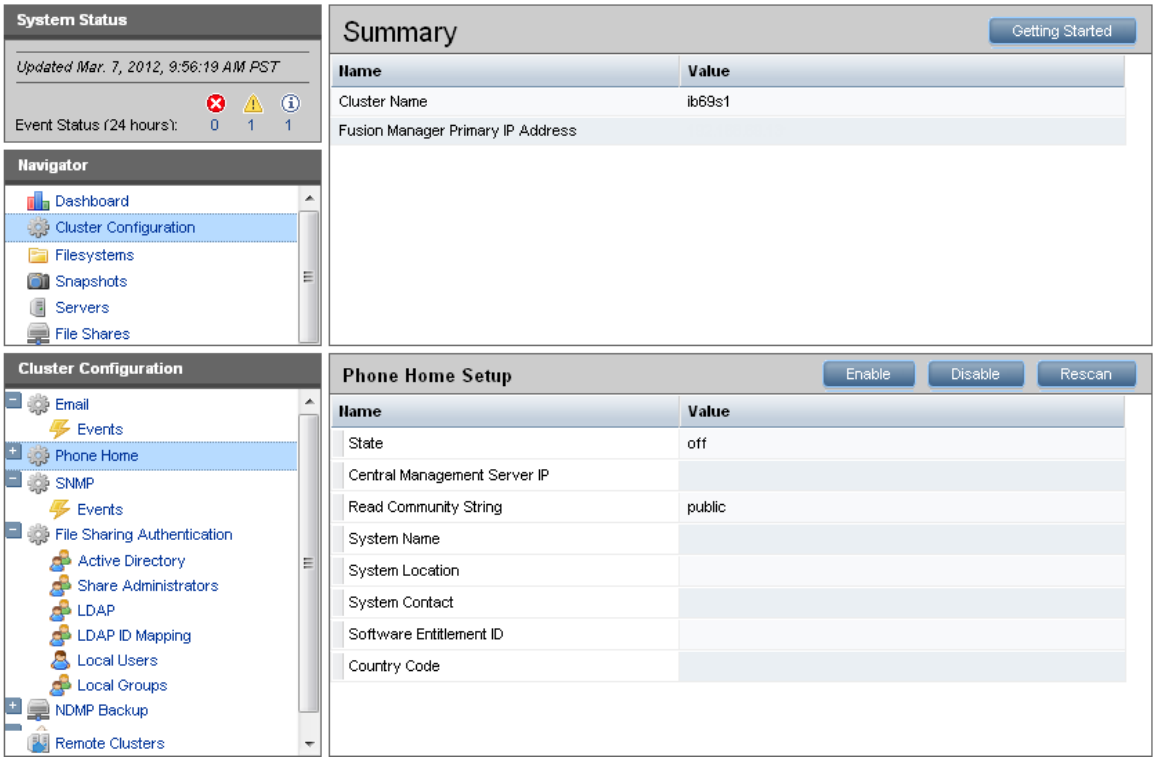

Click **Enable** to configure the settings on the Phone Home Settings dialog box. Skip the Software Entitlement ID field; it is not currently used.

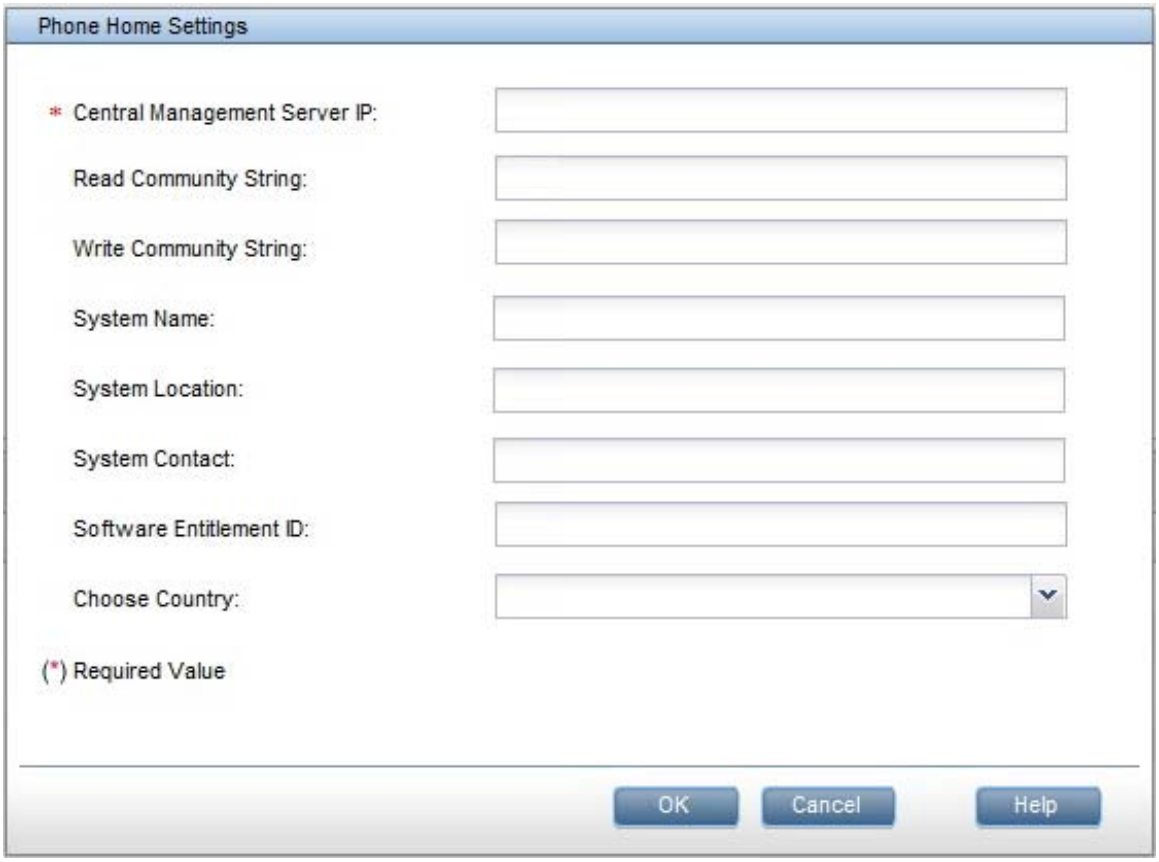

The time required to enable Phone Home depends on the number of devices in the cluster, with larger clusters requiring more time.

To configure Phone Home settings from the CLI, use the following command:

```
ibrix_phonehome -c -i <IP Address of the Central Management Server> -P
Country Name [-z Software Entitlement ID] [-r Read Community] [-w Write
Community] [-t System Contact] [-n System Name] [-o System Location]
```

For example:

```
ibrix_phonehome -c -i 99.2.4.75 -P US -r public -w private -t Admin -n
SYS01.US -o Colorado
```

Next, configure Insight Remote Support for the version of HP SIM you are using:

- HP SIM 7.1 and IRS 5.7. See "Configuring Insight Remote Support for HP SIM 7.1 and IRS 5.7" (page 39).
- HP SIM 6.3 and IRS 5.6. See "Configuring Insight Remote Support for HP SIM 6.3 and IRS 5.6" (page 42).

Configuring Insight Remote Support for HP SIM 7.1 and IRS 5.7

To configure Insight Remote Support, complete these steps:

1. Configure Entitlements for the servers and chassis in your system.
2. Discover devices on HP SIM.

Configuring Entitlements for servers and storage

Expand Phone Home in the lower Navigator. When you select **Servers**, or **Storage**, the GUI displays the current Entitlements for that type of device. The following example shows Entitlements for the servers in the cluster.

NOTE: The **Chassis** selection does not apply to 9300 or 9320 systems.

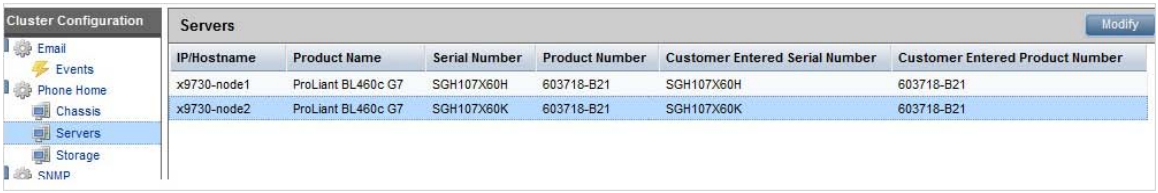

To configure Entitlements, select a device and click **Modify** to open the dialog box for that type of device. The following example shows the Server Entitlement dialog box. The customer-entered serial number and product number are used for warranty checks at HP Support.

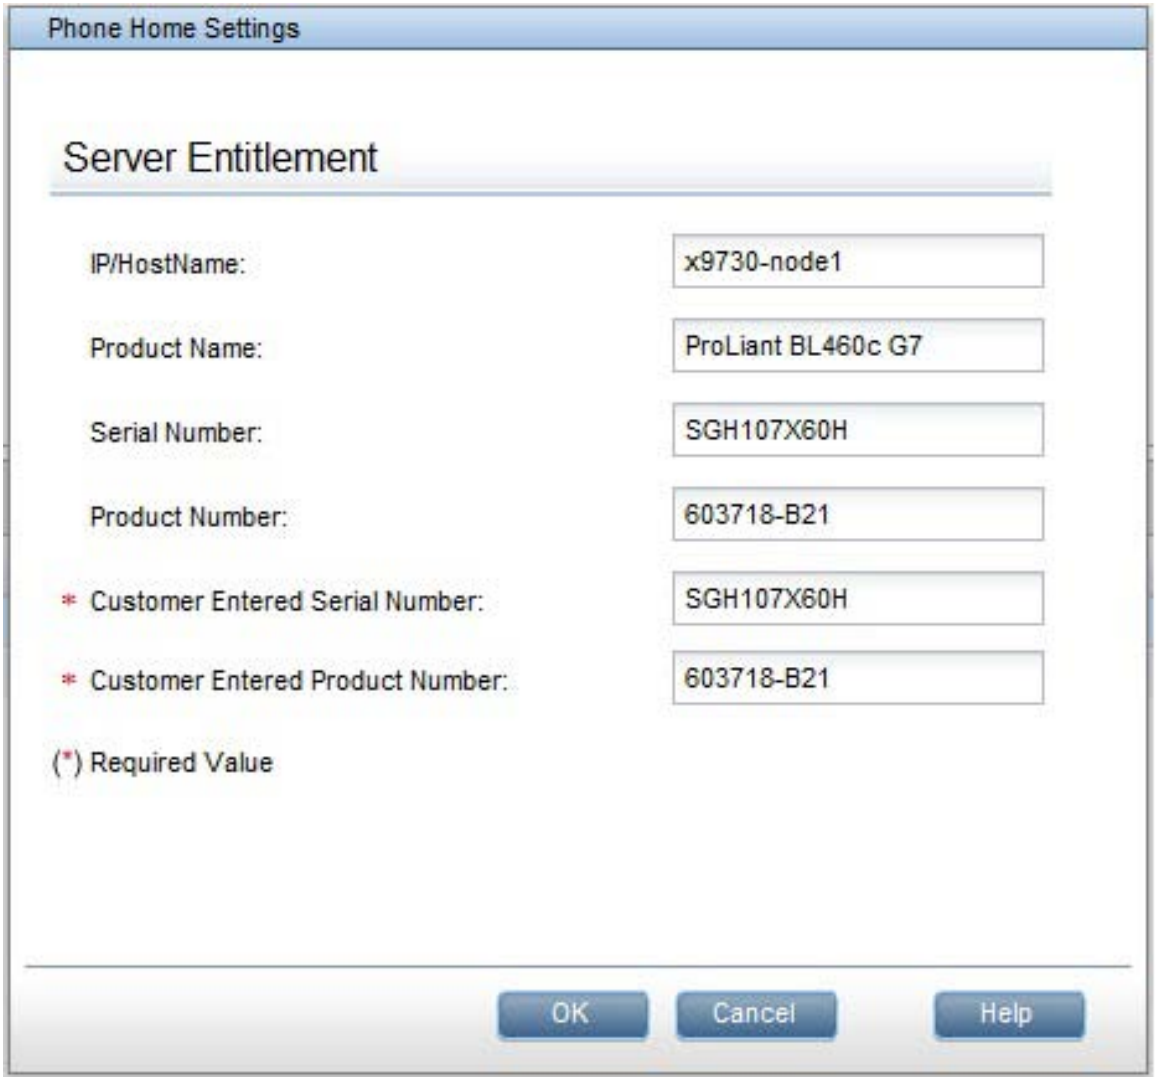

Use the following commands to entitle devices from the CLI. The commands must be run for each device present in the cluster.

Entitle a server:

`ibrix_phonehome -e -h <Host Name> -b <Customer Entered Serial Number> -g <Customer Entered Product Number>`

Enter the *Host Name* parameter exactly as it is listed by the `ibrix_fm -l` command.

Entitle storage (MSA):

`ibrix_phonehome -e -i <Management IP Address of the Storage> -b <Customer Entered Serial Number> -g <Customer Entered Product Number>`

Device discovery

HP Systems Insight Manager (SIM) uses the SNMP protocol to discover and identify StoreAll systems automatically. On HP SIM, open **Options** > **Discovery** > **New**. Select **Discover a group of systems**, and then enter the discovery name and the Fusion Manager IP address on the New Discovery dialog box.

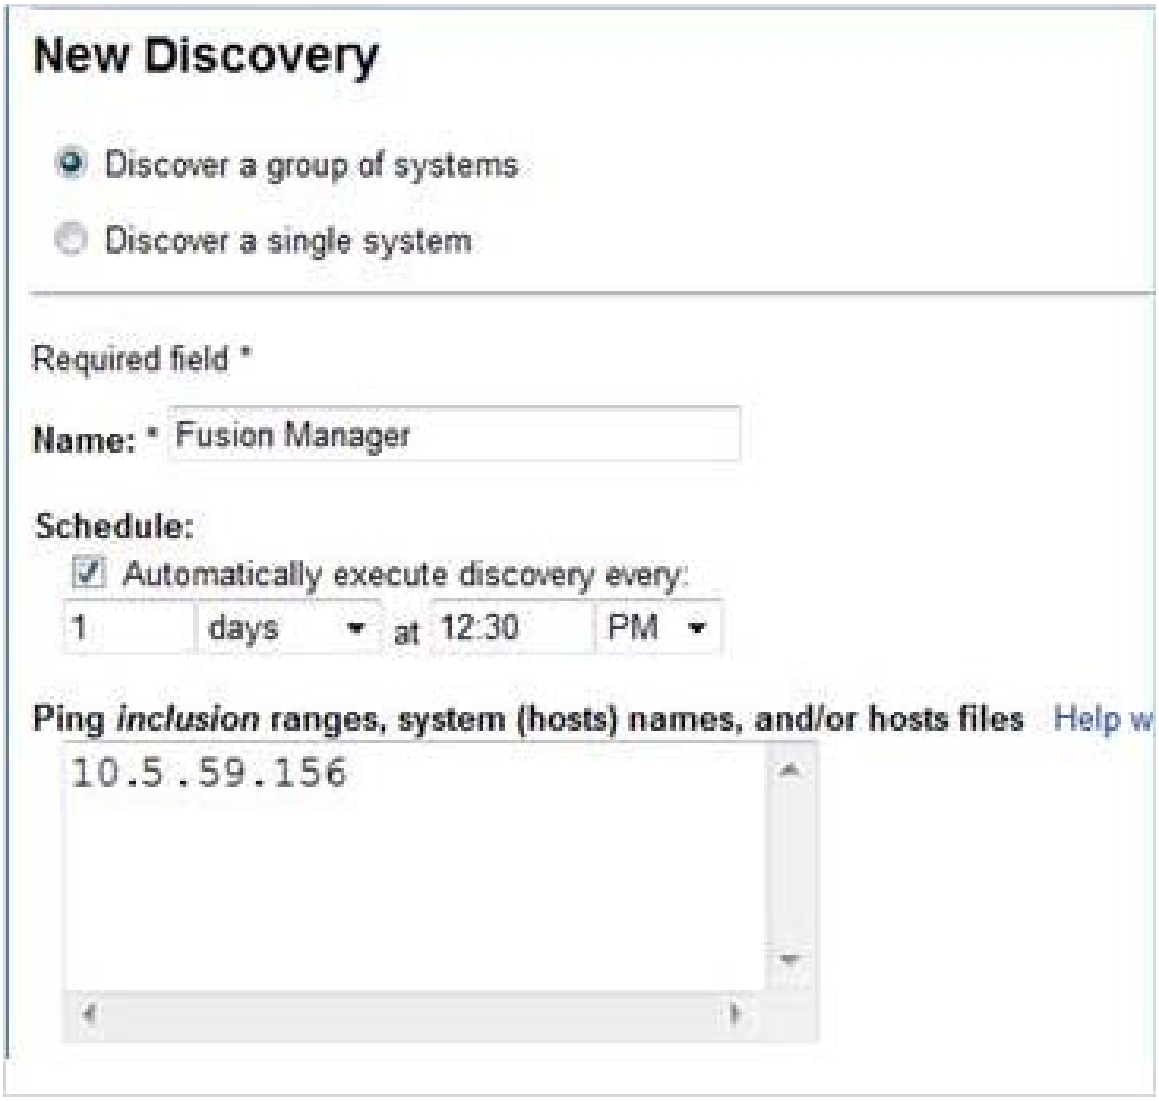

Enter the read community string on the **Credentials** > **SNMP** tab. This string should match the Phone Home read community string. If the strings are not identical, the Fusion Manager IP might be discovered as "Unknown."

Devices are discovered as described in the following table.

Device	Discovered as	
Fusion Manager IP	System Type: System Subtype: Product Model:	Fusion Manager 9000 HP 9000 Solution
File serving nodes	System Type: System Subtype: Product Model:	Storage Device 9000, Storage, HP ProLiant HP 9320 NetStor FSN(ProLiant DL380 G7) HP 9320 NetStor FSN(ProLiant DL380 G6) HP 9300 NetStor FSN(ProLiant DL380 G7) HP 9300 NetStor FSN(ProLiant DL380 G6)

The following example shows discovered devices on HP SIM 7.1.

File serving nodes and MSA arrays are associated with the Fusion Manager IP address. In HP SIM, select **Fusion Manager** and open the Systems tab. Then select **Associations** to view the devices.

You can view all StoreAll devices under **Systems by Type** > **Storage System** > **Scalable Storage Solutions** > **All 9000 Systems**

Configuring Insight Remote Support for HP SIM 6.3 and IRS 5.6

Discovering devices in HP SIM

HP Systems Insight Manager (SIM) uses the SNMP protocol to discover and identify StoreAll systems automatically. On HP SIM, open **Options** > **Discovery** > **New**, and then select **Discover a group of systems**. On the New Discovery dialog box, enter the discovery name and the IP addresses of the devices to be monitored. For more information, see the HP SIM 6.3 documentation.

NOTE: Each device in the cluster should be discovered separately.

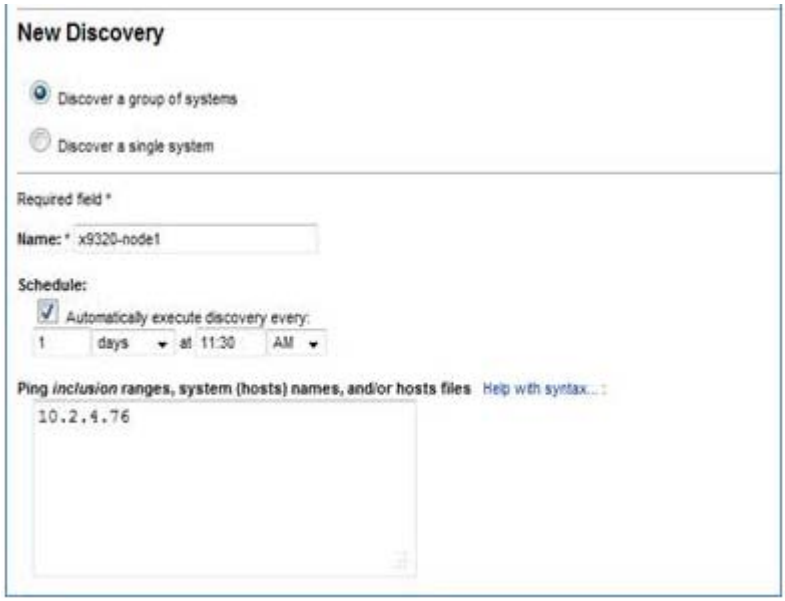
New Discovery
Discover a group of systems
Discover a single system
Required field *
Name: * x9320-node1
Schedule:
Automatically execute discovery every:
1 days at 11:30 AM
Ping inclusion ranges, system (hosts) names, and/or hosts files Help with syntax...
10.2.4.76

Enter the read community string on the **Credentials** > **SNMP** tab. This string should match the Phone Home read community string. If the strings are not identical, the device will be discovered as "Unknown."

The following example shows discovered devices on HP SIM 6.3. File serving nodes are discovered as ProLiant server.

HS Summary: 0 Critical 5 Major 0 Minor 2 Normal 0 Disabled 0 Unknown Total: 7
HS MP SW CW ES System Name System Type System Address Product Name
10.2.4.20 Server 10.2.59.104 ProLiant DL380 G6
10.2.4.30 Server 10.2.4.30 ProLiant DL380 G6
10.2.4.54 in Server 10.2.4.20 Management Processor 10.2.4.54
10.2.5.30 in Server 10.2.4.30 Management Processor 10.2.5.30 Integrated Lights-Out ...
win-79aqn41rpv8 Server 10.2.4.74

Configuring device Entitlements

Configure the CMS software to enable remote support for StoreAll systems. For more information, see "Using the Remote Support Setting Tab to Update Your Client and CMS Information" and "Adding Individual Managed Systems" in the *HP Insight Remote Support Advanced A.05.50 Operations Guide*.

Enter the following custom field settings in HP SIM:

- **Custom field settings for 9300/9320**

  Servers are discovered with their IP addresses. When a server is discovered, edit the system properties on the HP Systems Insight Manager. Locate the Entitlement Information section of the Contract and Warranty Information page and update the following:

  - Enter the StoreAll enclosure product number as the Customer-Entered product number
  - Enter `9000` as the Custom Delivery ID
  - Select the System Country Code
  - Enter the appropriate Customer Contact and Site Information details

- **Custom field settings for MSA Storage Management Utility**

  Configure SNMP settings on the MSA Storage Management Utility. (For more information, see "Configuring SNMP event notification in SMU" in the *2300 Modular Smart Array Reference Guide* This document is available at http://www.hp.com/support/manuals. On the Manuals page, select **storage** >**Disk Storage Systems** > **MSA Disk Arrays** >**HP 2000sa G2 Modular Smart Array** or **HP P2000 G3 MSA Array Systems**.) Refer to the *HP StorageWorks 2xxx Modular Smart Array Reference Guide* for other MSA versions.

  A Modular Storage Array (MSA) unit should be discovered with its IP address. Once discovered, locate the Entitlement Information section of the Contract and Warranty Information page and update the following:

  - Enter `9000` as the Custom Delivery ID
  - Select the System Country Code
  - Enter the appropriate Customer Contact and Site Information details

- **Contract and Warranty Information**

  Under Entitlement Information, specify the Customer-Entered serial number, Customer-Entered product number, System Country code, and Custom Delivery ID.

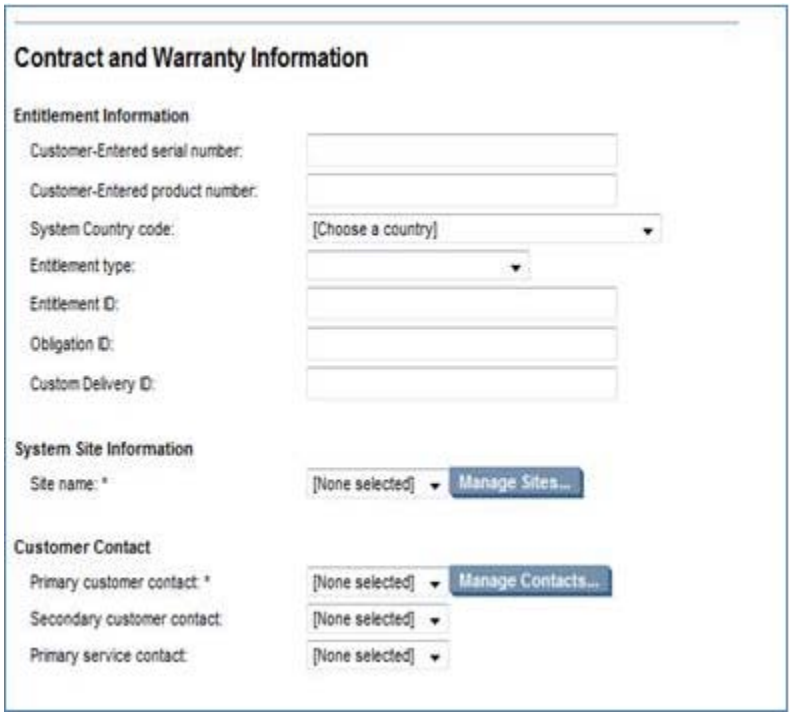

NOTE: For storage support on 9300 systems, do not set the Custom Delivery ID. (The MSA is an exception; the Custom Delivery ID is set as previously described.)

Verifying device entitlements

To verify the entitlement information in HP SIM, complete the following steps:

1. Go to **Remote Support Configuration and Service**s and select the **Entitlement** tab.
2. Check the devices discovered.

   NOTE: If the system discovered on HP SIM does not appear on the Entitlement tab, click **Synchronize RSE**.

3. Select **Entitle Checked** from the Action List.
4. Click **Run Action**.
5. When the entitlement check is complete, click **Refresh**.

   NOTE: If the system discovered on HP SIM does not appear on the Entitlement tab, click **Synchronize RSE**.

The devices you entitled should be displayed as green in the ENT column on the Remote Support System List dialog box.

If a device is red, verify that the customer-entered serial number and part number are correct and then rediscover the devices.

Testing the Insight Remote Support configuration

To determine whether the traps are working properly, send a generic test trap with the following command:

```
snmptrap -v1 -c public <CMS IP> .1.3.6.1.4.1.232 <Managed System IP> 6
11003 1234 .1.3.6.1.2.1.1.5.0 s test .1.3.6.1.4.1.232.11.2.11.1.0 i 0
.1.3.6.1.4.1.232.11.2.8.1.0 s "IBRIX remote support testing"
```

For example, if the CMS IP address is 99.2.2.2 and the StoreAll node is 99.2.2.10, enter the following:

```
snmptrap -v1 -c public 99.2.2.2 .1.3.6.1.4.1.232 99.2.2.10 6 11003 1234
.1.3.6.1.2.1.1.5.0 s test .1.3.6.1.4.1.232.11.2.11.1.0 i 0
.1.3.6.1.4.1.232.11.2.8.1.0 s "IBRIX remote support testing"
```

Updating the Phone Home configuration

The Phone Home configuration should be synchronized after you add or remove devices in the cluster. The operation enables Phone Home on newly added devices (servers, storage, and chassis) and removes details for devices that are no longer in the cluster. On the GUI, select **Cluster**

Configuration in the upper Navigator, select **Phone Home** in the lower Navigator, and click **Rescan** on the Phone Home Setup panel.

On the CLI, run the following command:

```
ibrix_phonehome -s
```

Disabling Phone Home

When Phone Home is disabled, all Phone Home information is removed from the cluster and hardware and software are no longer monitored. To disable Phone Home on the GUI, click **Disable** on the Phone Home Setup panel. On the CLI, run the following command:

```
ibrix_phonehome -d
```

Troubleshooting Insight Remote Support

Devices are not discovered on HP SIM

Verify that cluster networks and devices can access the CMS. Devices will not be discovered properly if they cannot access the CMS.

The maximum number of SNMP trap hosts has already been configured

If this error is reported when you configure Phone Home, the maximum number of trapsink IP addresses have already been configured. For MSA devices, the maximum number of trapsink IP addresses is 3. Manually remove a trapsink IP address from the device and then rerun the Phone Home configuration to allow Phone Home to add the CMS IP address as a trapsink IP address.

A cluster node was not configured in Phone Home

If a cluster node was down during the Phone Home configuration, the log file will include the following message:

```
SEVERE: Sent event server.status.down: Server <server name> down
```

When the node is up, rescan Phone Home to add the node to the configuration. See "Updating the Phone Home configuration" (page 45).

Fusion Manager IP is discovered as "Unknown"

Verify that the read community string entered in HP SIM matches the Phone Home read community string.

Also run `snmpwalk` on the VIF IP and verify the information:

```
# snmpwalk -v 1 -c <read community string> <FM VIF IP> .1.3.6.1.4.1.18997
```

Discovered device is reported as unknown on CMS

Run the following command on the file serving node to determine whether the Insight Remote Support services are running:

```
# service snmpd status
# service hpsmhd status
# service hp-snmp-agents status
```

If the services are not running, start them:

```
# service snmpd start
# service hpsmhd start
# service hp-snmp-agents start
```

Alerts are not reaching the CMS

If nodes are configured and the system is discovered properly but alerts are not reaching the CMS, verify that a `trapif` entry exists in the `cma.conf` configuration file on the file serving nodes.

Device Entitlement tab does not show GREEN

If the Entitlement tab does not show GREEN, verify the Customer-Entered serial number and part number or the device.

SIM Discovery

On SIM discovery, use the option **Discover a Group of Systems** for any device discovery.

4 Configuring virtual interfaces for client access

StoreAll software uses a cluster network interface to carry Fusion Manager traffic and traffic between file serving nodes. This network is configured as `bond0` when the cluster is installed. To provide failover support for the Fusion Manager, a virtual interface is created for the cluster network interface.

Although the cluster network interface can carry traffic between file serving nodes and clients, HP recommends that you configure one or more user network interfaces for this purpose.

To provide high availability for a user network, you should configure a bonded virtual interface (VIF) for the network and then set up failover for the VIF. This method prevents interruptions to client traffic. If necessary, the file serving node hosting the VIF can fail over to its backup server, and clients can continue to access the file system through the backup server.

StoreAll systems also support the use of VLAN tagging on the cluster and user networks. See "Configuring VLAN tagging" (page 51) for an example.

Network and VIF guidelines

To provide high availability, the user interfaces used for client access should be configured as bonded virtual interfaces (VIFs). Note the following:

- Nodes needing to communicate for file system coverage or for failover must be on the same network interface. Also, nodes set up as a failover pair must be connected to the same network interface.
- Use a Gigabit Ethernet port (or faster) for user networks.
- NFS, SMB, FTP, and HTTP clients can use the same user VIF. The servers providing the VIF should be configured in backup pairs, and the NICs on those servers should also be configured for failover. See "Configuring High Availability on the cluster" in the administrator guide for information about performing this configuration from the GUI.
- For Linux and Windows StoreAll clients, the servers hosting the VIF should be configured in backup pairs. However, StoreAll clients do not support backup NICs. Instead, StoreAll clients should connect to the parent bond of the user VIF or to a different VIF.
- Ensure that your parent bonds, for example bond0, have a defined route:
  1. Check for the default Linux OS route/gateway for each parent interface/bond that was defined during the HP StoreAll installation by entering the following command at the command prompt:

     ```
     # route
     ```

     The output from the command is the following:

     ```
     Destination     Gateway         Genmask         Flags   Metric  Ref  Use  Iface
     15.226.48.0       *           255.255.255.0     U        0       0    0   bond1
     172.16.0.0        *           255.255.248.0     U        0       0    0   bond0
     169.254.0.0       *           255.255.0.0       U        0       0    0   bond1
     default         15.226.48.1   0.0.0.0           UG       0       0    0   bond1
     Kernel IP routing table
     ```

     The default destination is the default gateway/route for Linux. The default destination, which was defined during the HP StoreAll installation, had the operating system default gateway defined but not for StoreAll.
  2. Display network interfaces controlled by StoreAll by entering the following command at the command prompt:

     ```
     # ibrix_nic -l
     ```

     Notice if the "ROUTE" column is unpopulated for IFNAME.

```
[root@ib50-80 ~]# ibrix_nic -l
HOST      IFNAME  TYPE     STATE          IP_ADDRESS     MAC_ADDRESS BACKUP_HOST BACKUP_IF  ROUTE       VLAN_TAG LINKMON
--------  ------  -------  -------------  -------------  ----------- ----------- ----------  ----------  -------- --------
ib50-81 bond0     Cluster  Up, LinkUp     172.16.0.81    00:00:00:00:11                      172.16.0.254            No
ib50-81 bond0:1   User     Up, LinkUp     172.16.0.181   00:00:00:00:11 ib50-82     bond0:2                      No
ib50-81 bond0:2   User                                   00:00:00:00:11                                          No
ib50-82 bond0     Cluster  Up, LinkUp     172.16.0.82    00:00:00:00:12                      172.16.0.254        No
ib50-82 bond0:1   User     Up, LinkUp     172.16.0.182   00:00:00:00:12 ib50-81     bond0:2                      No
ib50-82 bond0:2   User                                   00:00:00:00:12                                          No
ib50-81 [Active FM Nonedit] bond0:0 Cluster   Up, LinkUp   (ActiveFM) 172.16.0.281   No
```

3. To assign the IFNAME a default route for the parent cluster bond and the user VIFS assigned to FSNs for use with SMB/NFS, enter the following `ibrix_nic` command at the command prompt:

   ```
   # ibrix_nic -r -n IFNAME -h HOSTNAME-A -R <ROUTE_IP>
   ```

4. Configure backup monitoring, as described in "Configuring backup servers" (page 49).

Creating a bonded VIF

NOTE: The examples in this chapter use the unified network and create a bonded VIF on `bond0`. If your cluster uses a different network layout, create the bonded VIF on a user network bond such as `bond1`.

Use the following procedure to create a bonded VIF (`bond0:1` in this example):

1. If high availability (automated failover) is configured on the servers, disable it. Run the following command on the Fusion Manager:

   ```
   # ibrix_server -m -U
   ```

2. Identify the `bond0:1` VIF:

   ```
   # ibrix_nic -a -n bond0:1 -h node1,node2,node3,node4
   ```

3. Assign an IP address to the `bond1:1` VIFs on each node. In the command, `-I` specifies the IP address, `-M` specifies the netmask, and `-B` specifies the broadcast address:

   ```
   # ibrix_nic -c -n bond0:1 -h node1 -I 16.123.200.201 -M 255.255.255.0 -B 16.123.200.255
   # ibrix_nic -c -n bond0:1 -h node2 -I 16.123.200.202 -M 255.255.255.0 -B 16.123.200.255
   # ibrix_nic -c -n bond0:1 -h node3 -I 16.123.200.203 -M 255.255.255.0 -B 16.123.200.255
   # ibrix_nic -c -n bond0:1 -h node4 -I 16.123.200.204 -M 255.255.255.0 -B 16.123.200.255
   ```

Configuring backup servers

The servers in the cluster are configured in backup pairs. If this step was not done when your cluster was installed, assign backup servers for the `bond0:1` interface. In the following example, `node1` is the backup for `node2`, `node2` is the backup for `node1`, `node3` is the backup for `node4`, and `node4` is the backup for `node3`.

1. Add the VIF:

   ```
   # ibrix_nic -a -n bond0:2 -h node1,node2,node3,node4
   ```

2. Set up a backup server for each VIF:

   ```
   # ibrix_nic -b -H node1/bond0:1,node2/bond0:2
   # ibrix_nic -b -H node2/bond0:1,node1/bond0:2
   # ibrix_nic -b -H node3/bond0:1,node4/bond0:2
   # ibrix_nic -b -H node4/bond0:1,node3/bond0:2
   ```

Configuring NIC failover

NIC monitoring should be configured on VIFs that will be used by NFS, SMB, FTP, or HTTP.

IMPORTANT: When configuring NIC monitoring, use the same backup pairs that you used when configuring standby servers.

For example:

```
# ibric_nic -m -h node1 -A node2/bond0:1
# ibric_nic -m -h node2 -A node1/bond0:1
# ibric_nic -m -h node3 -A node4/bond0:1
# ibric_nic -m -h node4 -A node3/bond0:1
```

Configuring automated failover

To enable automated failover for your file serving nodes, execute the following command:

```
ibrix_server -m [-h SERVERNAME]
```

Example configuration

This example uses two nodes, `ib50-81` and `ib50-82`. These nodes are backups for each other, forming a backup pair.

```
[root@ib50-80 ~]# ibrix_server -l
Segment Servers
===============
SERVER_NAME  BACKUP    STATE          HA   ID                                     GROUP
-----------  -------   ------------   ---  -------------------------------------  -----
ib50-81      ib50-82   Up             on   132cf61a-d25b-40f8-890e-e97363ae0d0b   servers
ib50-82      ib50-81   Up             on   7d258451-4455-484d-bf80-75c94d17121d   servers
```

All VIFs on `ib50-81` have backup (standby) VIFs on `ib50-82`. Similarly, all VIFs on `ib50-82` have backup (standby) VIFs on `ib50-81`. NFS, SMB, FTP, and HTTP clients can connect to `bond0:1` on either host. If necessary, the selected server will fail over to `bond0:2` on the opposite host. StoreAll clients could connect to `bond1` on either host, as these clients do not support or require NIC failover. (The following sample output shows only the relevant fields.)

```
[root@ib50-80 ~]# ibrix_nic -l
HOST     IFNAME      TYPE     STATE          IP_ADDRESS     MAC_ADDRESS       BACKUP_HOST  BACKUP_IF   ROUTE        VLAN_TAG  LINKMON
-------  ------      -------  -------------  -------------  ----------------  -----------  ----------  -----------  --------  -------
ib50-81  bond0       Cluster  Up, LinkUp     172.16.0.81    00:00:00:00:11                             172.16.0.254           No
ib50-81  bond0:1     User     Up, LinkUp     172.16.0.181   00:00:00:00:11    ib50-82      bond0:2                           No
ib50-81  bond0:2     User                    00:00:00:00:11                                                                  No
ib50-82  bond0       Cluster  Up, LinkUp     172.16.0.82    00:00:00:00:12                             172.16.0.254           No
ib50-82  bond0:1     User     Up, LinkUp     172.16.0.182   00:00:00:00:12    ib50-81      bond0:2                           No
ib50-82  bond0:2     User                    00:00:00:00:12
ib50-81  [Active FM Nonedit]  bond0:0  Cluster  Up, LinkUp  (ActiveFM)  172.16.0.281  No
```

Specifying VIFs in the client configuration

When you configure your clients, you may need to specify the VIF that should be used for client access.

NFS/SMB. Specify the VIF IP address of the servers (for example, `bond0:1`) to establish connection. You can also configure DNS round robin to ensure NFS or SMB client-to-server distribution. In both cases, the NFS/SMB clients will cache the initial IP they used to connect to the respective share, usually until the next reboot.

FTP. When you add an FTP share on the Add FTP Shares dialog box or with the `ibrix_ftpshare` command, specify the VIF as the IP address that clients should use to access the share.

HTTP. When you create a virtual host on the Create Vhost dialog box or with the `ibrix_httpvhost` command, specify the VIF as the IP address that clients should use to access shares associated with the Vhost.

StoreAll clients. Use the following command to prefer the appropriate user network. Execute the command once for each destination host that the client should contact using the specified interface.

```
ibrix_client -n -h SRCHOST -A DESTNOST/IFNAME
```

For example:

```
ibrix_client -n -h client12.mycompany.com -A ib50-81.mycompany.com/bond1
```

NOTE: Because the backup NIC cannot be used as a preferred network interface for StoreAll clients, add one or more user network interfaces to ensure that HA and client communication work together.

Configuring VLAN tagging

VLAN capabilities provide hardware support for running multiple logical networks over the same physical networking hardware. To allow multiple packets for different VLANs to traverse the same physical interface, each packet must have a field added that contains the VLAN tag. The tag is a small integer number that identifies the VLAN to which the packet belongs. When an intermediate switch receives a "tagged" packet, it can make the appropriate forwarding decisions based on the value of the tag.

When set up properly, StoreAll systems support VLAN tags being transferred all of the way to the file serving node network interfaces. The ability of file serving nodes to handle the VLAN tags natively in this manner makes it possible for the nodes to support multiple VLAN connections simultaneously over a single bonded interface.

Linux networking tools such as `ifconfig` display a network interface with an associated VLAN tag using a device label with the form `bond#.<VLAN_id>`. For example, if the first bond created by StoreAll has a VLAN tag of `30`, it will be labeled `bond0.30`.

It is also possible to add a VIF on top of an interface that has an associated VLAN tag. In this case, the device label of the interface takes the form `bond#.<VLAN_id>.<VVIF_label>`. For example, if a VIF with a label of `2` is added for the `bond0.30` interface, the new interface device label will be `bond0.30:2`.

The following commands show configuring a bonded VIF and backup nodes for a unified network topology using the `10.10.x.y` subnet. VLAN tagging is configured for hosts `ib142-129` and `ib142-131` on the `51` subnet.

Add the `bond0.51` interface with the VLAN tag:

```
# ibrix_nic -a -n bond0.51 -h ib142-129
# ibrix_nic -a -n bond0.51 -h ib142-131
```

Assign an IP address to the `bond0:51` VIFs on each node:

```
# ibrix_nic -c -n bond0.51 -h ib142-129 -I 192.168.51.101 -M 255.255.255.0
# ibrix_nic -c -n bond0.51 -h ib142-131 -I 192.168.51.102 -M 255.255.255.0
```

Add the `bond0.51:2` VIF on top of the interface:

```
# ibrix_nic -a -n bond0.51:2 -h ib142-131
# ibrix_nic -a -n bond0.51:2 -h ib142-129
```

Configure backup nodes:

```
# ibrix_nic -b -H ib142-129/bond0.51,ib142-131/bond0.51:2
# ibrix_nic -b -H ib142-131/bond0.51,ib142-129/bond0.51:2
```

Create the user FM VIF:

```
ibrix_fm -c 192.168.51.125 -d bond0.51:1 -n 255.255.255.0 -v user
```

For more information about VLAG tagging, see the *HP StoreAll Storage Network Best Practices Guide*.

Configuring link state monitoring for iSCSI network interfaces

Do not configure link state monitoring for user network interfaces or VIFs that will be used for SMB or NFS. Link state monitoring is supported only for use with iSCSI storage network interfaces, such as those provided with 9300 Gateway systems.

To configure link state monitoring on a 9300 system, use the following command:

```
ibrix_nic -N -h HOST -A IFNAME
```

To determine whether link state monitoring is enabled on an iSCSI interface, run the following command:

```
ibrix_nic -l
```

Next, check the LINKMON column in the output. The value `yes` means that link state monitoring is enabled; `no` means that it is not enabled.

5 Configuring failover

This chapter describes how to configure failover for agile management consoles, file serving nodes, network interfaces, and HBAs.

Agile management consoles

The agile Fusion Manager maintains the cluster configuration and provides graphical and command-line user interfaces for managing and monitoring the cluster. The agile Fusion Manager is installed on all file serving nodes when the cluster is installed. The Fusion Manager is active on one node, and is passive on the other nodes. This is called an *agile* Fusion Manager configuration.

Agile Fusion Manager modes

An agile Fusion Manager can be in one of the following modes:

- **active**. In this mode, the Fusion Manager controls console operations. All cluster administration and configuration commands must be run from the active Fusion Manager.
- **passive**. In this mode, the Fusion Manager monitors the health of the active Fusion Manager. If the active Fusion Manager fails, the a passive Fusion Manager is selected to become the active console.
- **nofmfailover**. In this mode, the Fusion Manager does not participate in console operations. Use this mode for operations such as manual failover of the active Fusion Manager, StoreAll software upgrades, and server blade replacements.

Changing the mode

Use the following command to move a Fusion Manager to passive or nofmfailover mode:

```
ibrix_fm -m passive | nofmfailover [-P] [-A | -h <FMLIST>]
```

If the Fusion Manager was previously the active console, StoreAll software will select a new active console. A Fusion Manager currently in active mode can be moved to either passive or nofmfailover mode. A Fusion Manager in nofmfailover mode can be moved only to passive mode.

With the exception of the local node running the active Fusion Manager, the `-A` option moves all instances of the Fusion Manager to the specified mode. The `-h` option moves the Fusion Manager instances in *<FMLIST>* to the specified mode.

Viewing information about Fusion Managers

To view mode information, use the following command:

```
ibrix_fm -i
```

NOTE: If the Fusion Manager was not installed in an agile configuration, the output will report `FusionServer: fusion manager name not set! (active, quorum is not configured)`.

When a Fusion Manager is installed, it is registered in the Fusion Manager configuration. To view a list of all registered management consoles, use the following command:

```
ibrix_fm -l
```

Agile Fusion Manager and failover

Using an agile Fusion Manager configuration provides high availability for Fusion Manager services. If the active Fusion Manager fails, the cluster virtual interface will go down. When the passive Fusion Manager detects that the cluster virtual interface is down, it will become the active

console. This Fusion Manager rebuilds the cluster virtual interface, starts Fusion Manager services locally, transitions into active mode, and take over Fusion Manager operation.

Failover of the active Fusion Manager affects the following features:

- **User networks**. The virtual interface used by clients will also fail over. Users may notice a brief reconnect while the newly active Fusion Manager takes over management of the virtual interface.
- **GUI**. You must reconnect to the Fusion Manager VIF after the failover.

Failing over the Fusion Manager manually

To fail over the active Fusion Manager manually, place the console into nofmfailover mode. Enter the following command on the node hosting the console:

```
ibrix_fm -m nofmfailover
```

The failover will take approximately one minute.

Run to see which node is now the active Fusion Manager, enter the following command:

```
ibrix_fm -i
```

The failed-over Fusion Manager remains in nofmfailover mode until it is moved to passive mode using the following command:

```
ibrix_fm -m passive
```

NOTE: A Fusion Manager cannot be moved from nofmfailover mode to active mode.

Configuring High Availability on the cluster

StoreAll High Availability provides monitoring for servers, NICs, and HBAs.

Server HA. Servers are configured in backup pairs, with each server in the pair acting as a backup for the other server. The servers in the backup pair must see the same storage. When a server is failed over, the ownership of its segments and its Fusion Manager services (if the server is hosting the active FM) move to the backup server.

NIC HA. When server HA is enabled, NIC HA provides additional triggers that cause a server to fail over to its backup server. For example, you can create a user VIF such as `bond0:2` to service SMB requests on a server and then designate the backup server as a standby NIC for `bond0:2`. If an issue occurs with `bond0:2` on a server, the server, including its segment ownership and FM services, will fail over to the backup server, and that server will now handle SMB requests going through `bond0:2`.

You can also fail over just the NIC to its standby NIC on the backup server.

HBA monitoring. This method protects server access to storage through an HBA. Most servers ship with an HBA that has two controllers, providing redundancy by design. Setting up StoreAll HBA monitoring is not commonly used for these servers. However, if a server has only a single HBA, you might want to monitor the HBA; then, if the server cannot see its storage because the single HBA goes offline or faults, the server and its segments will fail over.

You can set up automatic server failover and perform a manual failover if needed. If a server fails over, you must fail back the server manually.

When automatic HA is enabled, the Fusion Manager listens for heartbeat messages that the servers broadcast at one-minute intervals. The Fusion Manager initiates a server failover when it fails to receive five consecutive heartbeats. Failover conditions are detected more quickly when NIC HA is also enabled; server failover is initiated when the Fusion Manager receives a heartbeat message indicating that a monitored NIC might be down and the Fusion Manager cannot reach that NIC. If HBA monitoring is enabled, the Fusion Manager fails over the server when a heartbeat message indicates that a monitored HBA or pair of HBAs has failed.

What happens during a failover

The following actions occur when a server is failed over to its backup:

1. The Fusion Manager verifies that the backup server is powered on and accessible.
2. The Fusion Manager migrates ownership of the server's segments to the backup and notifies all servers and StoreAll clients about the migration. This is a persistent change. If the server is hosting the active FM, it transitions to another server.
3. If NIC monitoring is configured, the Fusion Manager activates the standby NIC and transfers the IP address (or VIF) to it.

Clients that were mounted on the failed-over server may experience a short service interruption while server failover takes place. Depending on the protocol in use, clients can continue operations after the failover or may need to remount the file system using the same VIF. In either case, clients will be unaware that they are now accessing the file system on a different server.

To determine the progress of a failover, view the Status tab on the GUI or execute the `ibrix_server -l` command. While the Fusion Manager is migrating segment ownership, the operational status of the node is `Up-InFailover` or `Down-InFailover`, depending on whether the node was powered up or down when failover was initiated. When failover is complete, the operational status changes to `Up-FailedOver` or `Down-FailedOver`. For more information about operational states, see "Monitoring the status of file serving nodes" (page 93).

Both automated and manual failovers trigger an event that is reported on the GUI.

Automated failover can be configured with the HA Wizard or from the command line.

Configuring automated failover with the HA Wizard

The HA wizard configures a backup server pair and, optionally, standby NICs on each server in the pair. It also configures a power source such as an iLO on each server. The Fusion Manager uses the power source to power down the server during a failover.

On the GUI, select **Servers** from the Navigator.

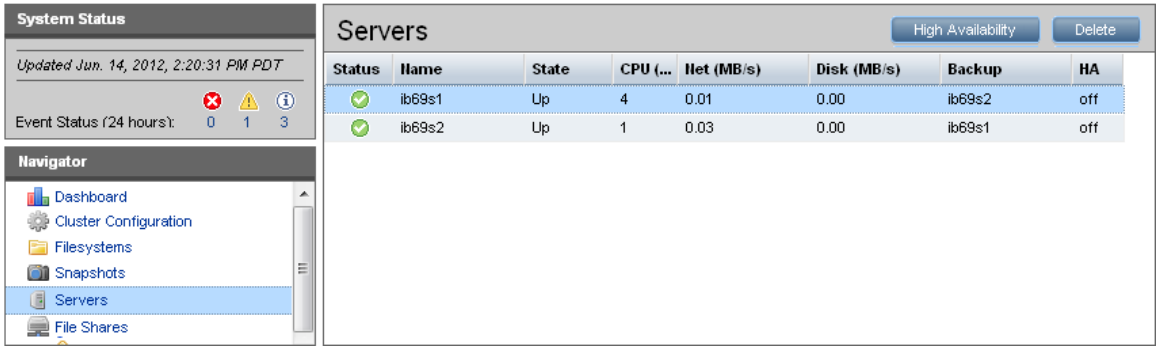

Click **High Availability** to start the wizard. Typically, backup servers are configured and server HA is enabled when your system is installed, and the Server HA Pair dialog box shows the backup pair configuration for the server selected on the Servers panel.

If necessary, you can configure the backup pair for the server. The wizard identifies the servers in the cluster that see the same storage as the selected server. Choose the appropriate server from the list.

The wizard also attempts to locate the IP addresses of the iLOs on each server. If it cannot locate an IP address, you will need to enter the address on the dialog box. When you have completed the information, click **Enable HA Monitoring and Auto-Failover for both servers**.

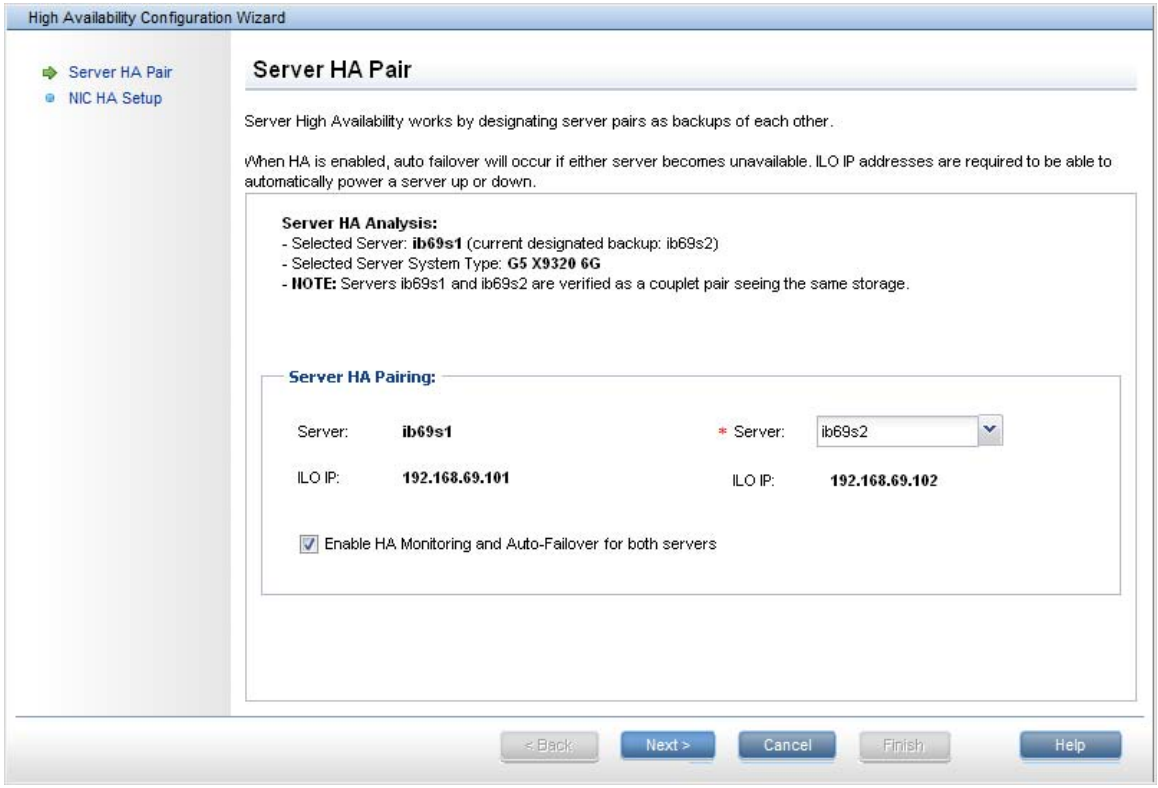
High Availability Configuration Wizard
Server HA Pair
NIC HA Setup
Server HA Pair
Server High Availability works by designating server pairs as backups of each other.
When HA is enabled, auto failover will occur if either server becomes unavailable. ILO IP addresses are required to be able to automatically power a server up or down.
Server HA Analysis:
- Selected Server: ib69s1 (current designated backup: ib69s2)
- Selected Server System Type: G5 X9320 6G
- NOTE: Servers ib69s1 and ib69s2 are verified as a couplet pair seeing the same storage.
Server HA Pairing:
Server: ib69s1
* Server: ib69s2
ILO IP: 192.168.69.101
ILO IP: 192.168.69.102
Enable HA Monitoring and Auto-Failover for both servers
< Back
Next >
Cancel
Finish
Help

Use the NIC HA Setup dialog box to configure NICs that will be used for data services such as SMB or NFS. You can also designate NIC HA pairs on the server and its backup and enable monitoring of these NICs.

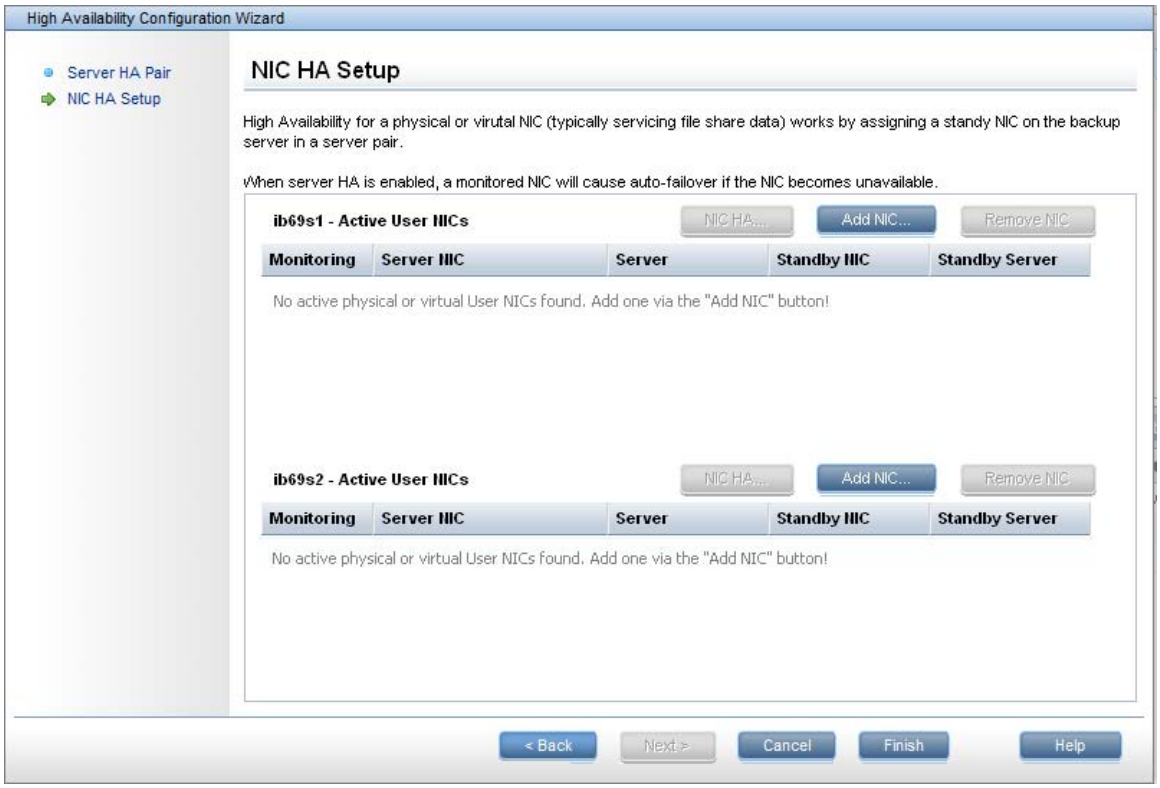
High Availability Configuration Wizard
Server HA Pair
NIC HA Setup
NIC HA Setup
High Availability for a physical or virutal NIC (typically servicing file share data) works by assigning a standy NIC on the backup server in a server pair.
When server HA is enabled, a monitored NIC will cause auto-failover if the NIC becomes unavailable.
ib69s1 - Active User NICs
NIC HA...
Add NIC...
Remove NIC
Monitoring
Server NIC
Server
Standby NIC
Standby Server
No active physical or virtual User NICs found. Add one via the "Add NIC" button!
ib69s2 - Active User NICs
NIC HA...
Add NIC...
Remove NIC
Monitoring
Server NIC
Server
Standby NIC
Standby Server
No active physical or virtual User NICs found. Add one via the "Add NIC" button!
< Back
Next >
Cancel
Finish
Help

For example, you can create a user VIF that clients will use to access an SMB share serviced by server ib69s1. The user VIF is based on an active physical network on that server. To do this, click **Add NIC** in the section of the dialog box for ib69s1.

On the Add NIC dialog box, enter a NIC name. In our example, the cluster uses the unified network and has only `bond0`, the active cluster FM/IP. We cannot use `bond0:0`, which is the management IP/VIF. We will create the VIF `bond0:1`, using `bond0` as the base. When you click **OK**, the user VIF is created.

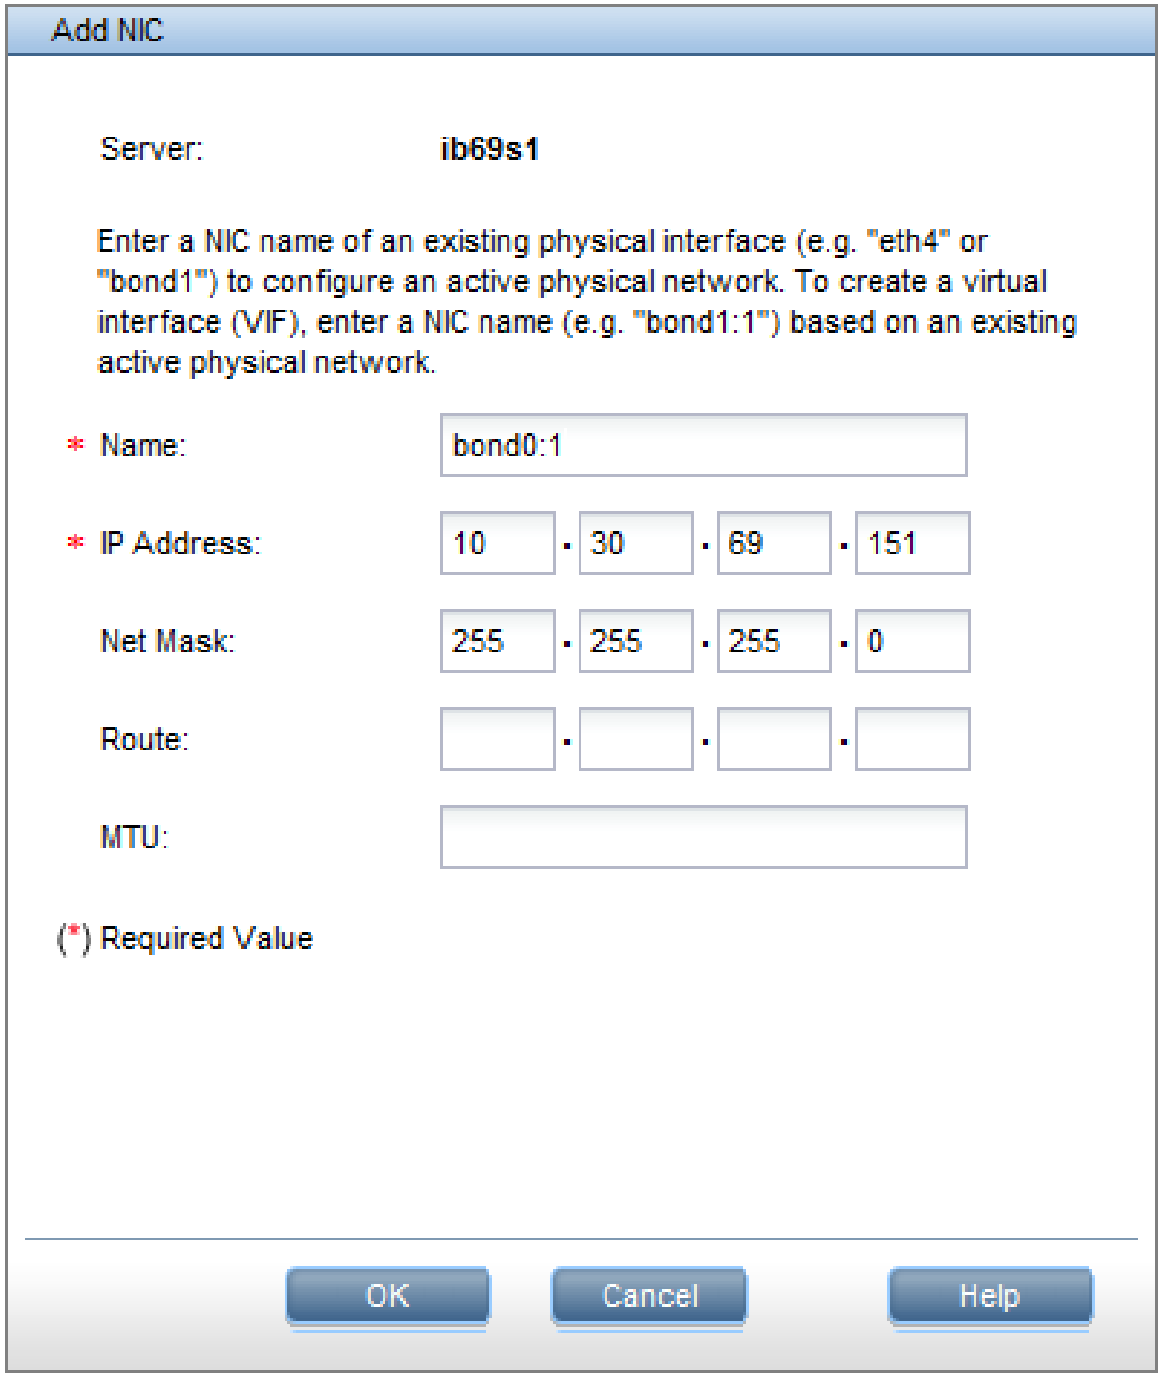
Add NIC
Server:
ib69s1
Enter a NIC name of an existing physical interface (e.g. "eth4" or "bond1") to configure an active physical network. To create a virtual interface (VIF), enter a NIC name (e.g. "bond1:1") based on an existing active physical network.
* Name:
bond0:1
* IP Address:
10 . 30 . 69 . 151
Net Mask:
255 . 255 . 255 . 0
Route:
MTU:
(*) Required Value
OK
Cancel
Help

The new, active user NIC appears on the NIC HA setup dialog box.

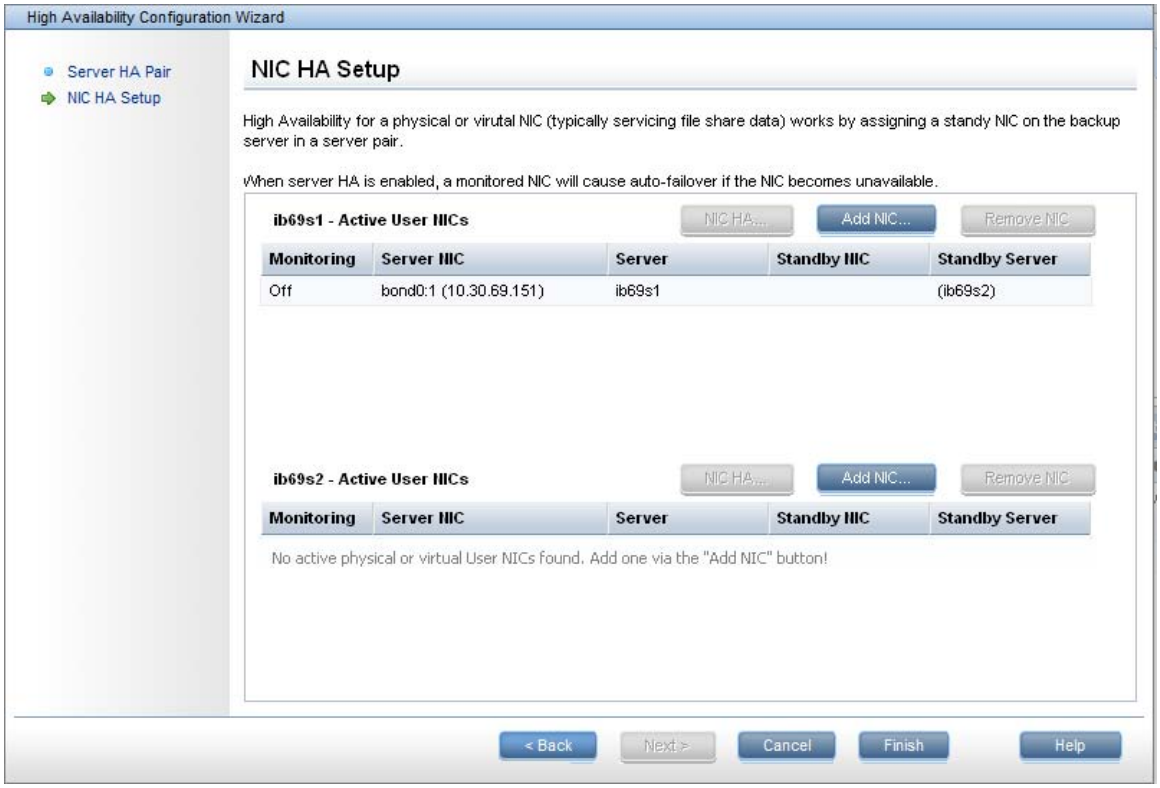

Next, enable NIC monitoring on the VIF. Select the new user NIC and click **NIC HA**. On the NIC HA Config dialog box, check **Enable NIC Monitoring**.

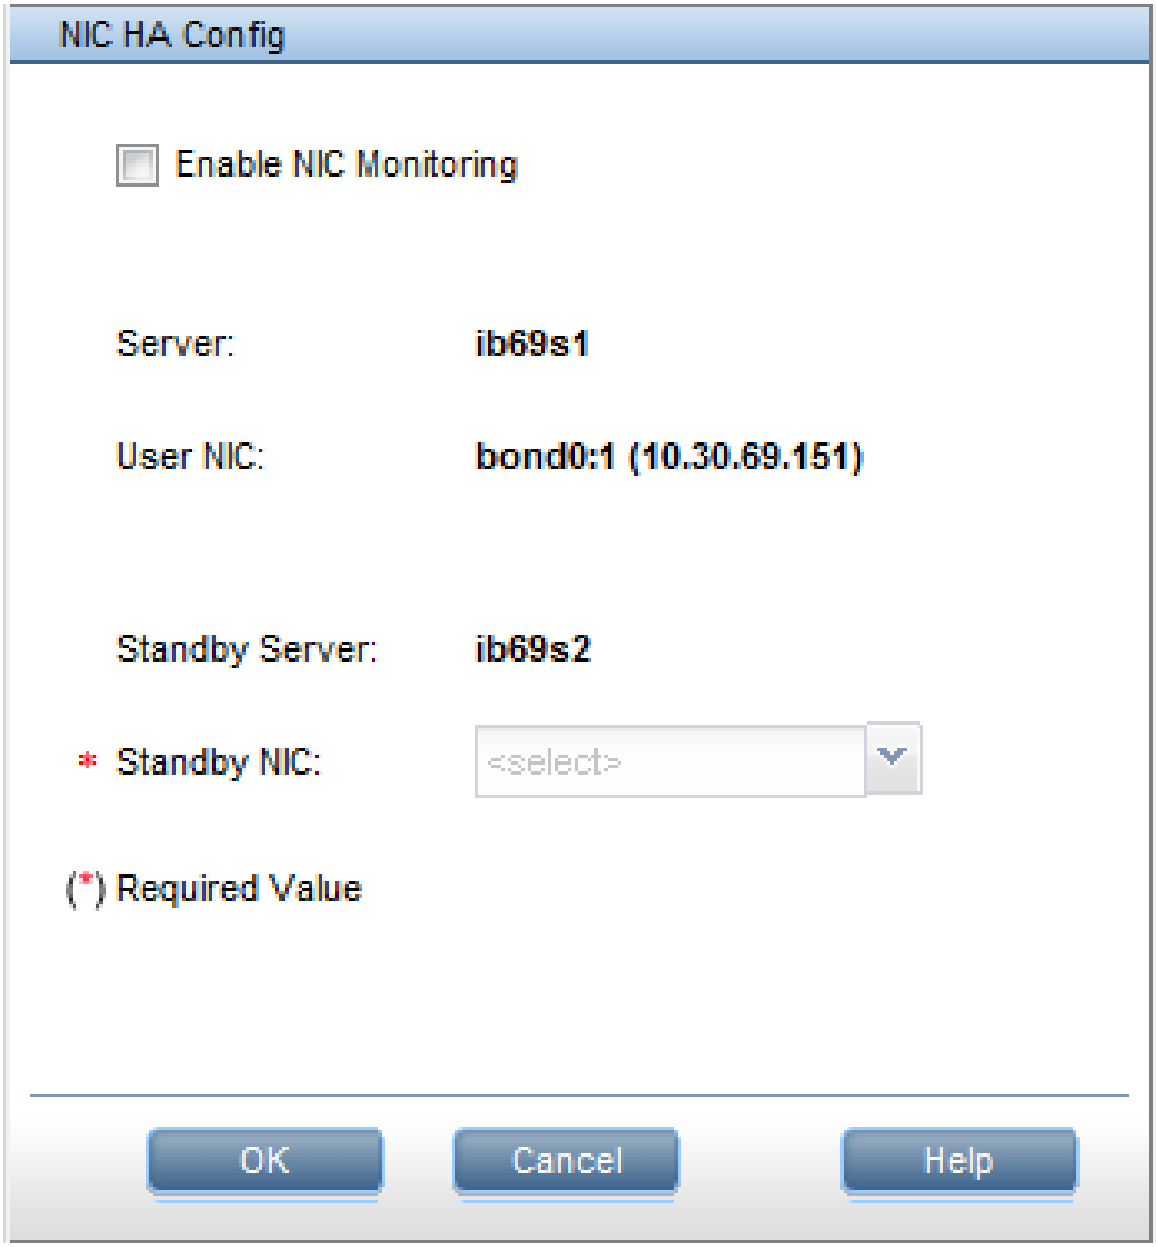

In the Standby NIC field, select **New Standby NIC** to create the standby on backup server ib69s2. The standby you specify must be available and valid. To keep the organization simple, we specified **`bond0:1`** as the **Name**; this matches the name assigned to the NIC on server ib69s1. When you click **OK**, the NIC HA configuration is complete.

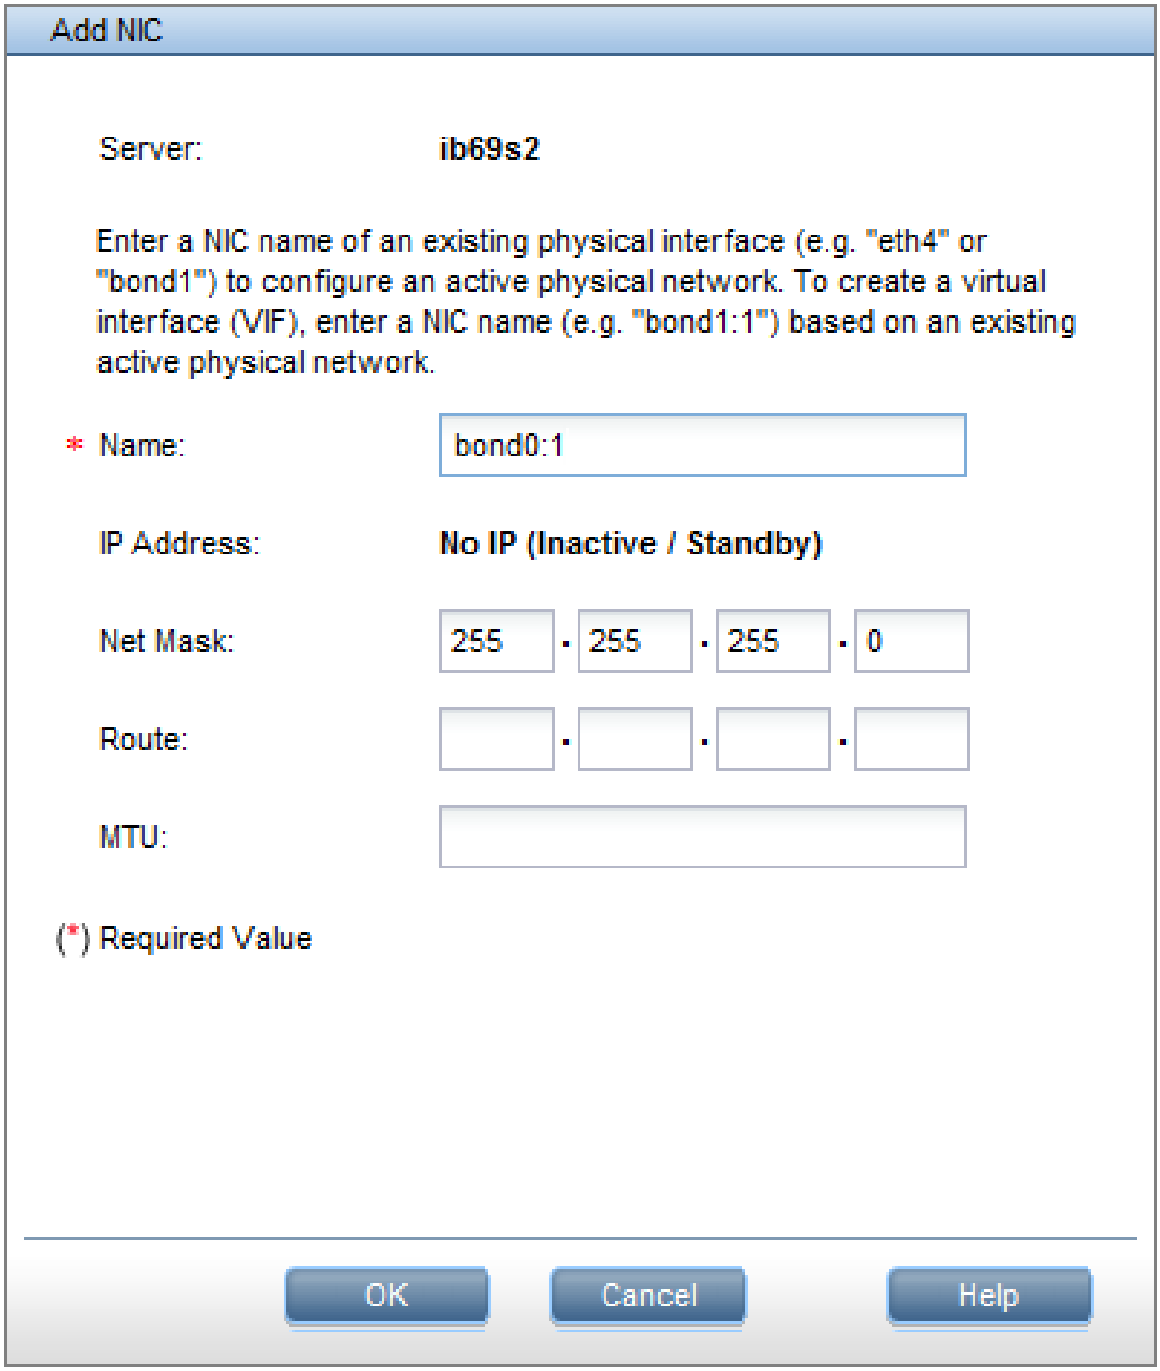
Add NIC
Server:
ib69s2
Enter a NIC name of an existing physical interface (e.g. "eth4" or "bond1") to configure an active physical network. To create a virtual interface (VIF), enter a NIC name (e.g. "bond1:1") based on an existing active physical network.
* Name:
bond0:1
IP Address:
No IP (Inactive / Standby)
Net Mask:
255 . 255 . 255 . 0
Route:
MTU:
(*) Required Value
OK
Cancel
Help

You can create additional user VIFs and assign standby NICs as needed. For example, you might want to add a user VIF for another share on server ib69s2 and assign a standby NIC on server ib69s1. You can also specify a physical interface such `eth4` and create a standby NIC on the backup server for it.

The NICs panel on the GUI shows the NICs on the selected server. In the following example, there are four NICs on server ib69s1: `bond0`, the active cluster FM/IP; `bond0:0`, the management IP/VIF (this server is hosting the active FM); `bond0:1`, the NIC created in this example; and `bond0:2`, a standby NIC for an active NIC on server ib69s2.

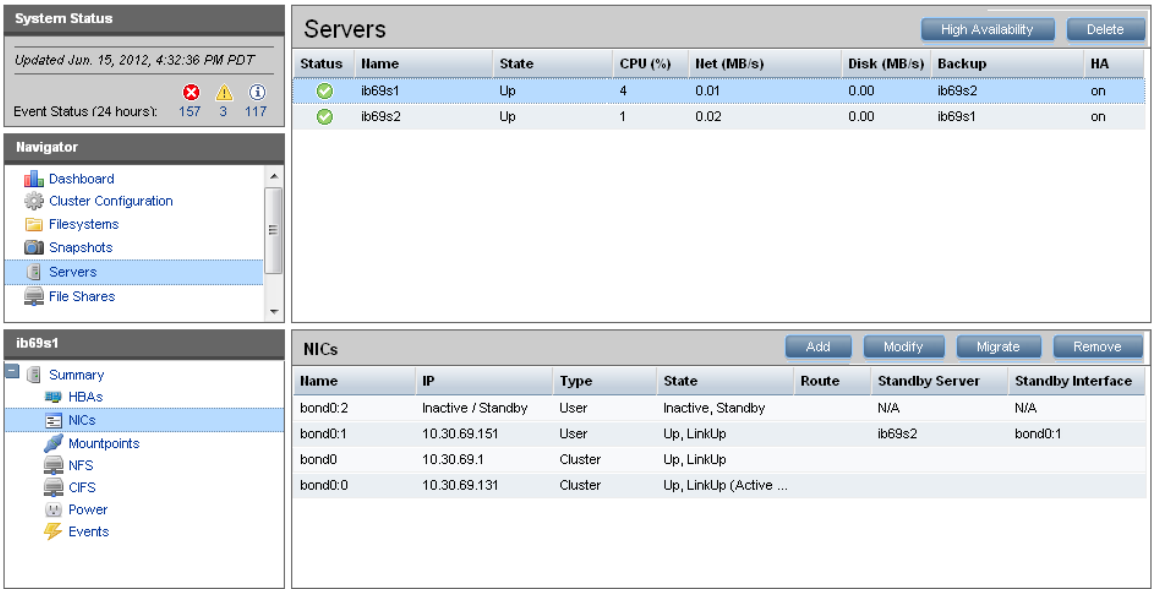

The NICs panel for the ib69s2, the backup server, shows that `bond0:1` is an inactive, standby NIC and `bond0:2` is an active NIC.

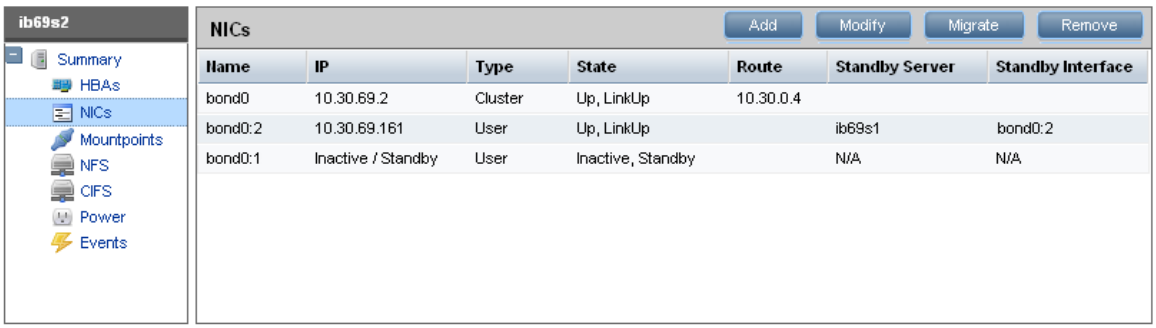

Changing the HA configuration

To change the configuration of a NIC, select the server on the Servers panel, and then select **NICs** from the lower Navigator. Click **Modify** on the NICs panel. The General tab on the Modify NIC Properties dialog box allows you change the IP address and other NIC properties. The NIC HA tab allows you to enable or disable HA monitoring and failover on the NIC and to change or remove the standby NIC. You can also enable link state monitoring if it is supported on your cluster. See "Configuring link state monitoring for iSCSI network interfaces" (page 51).

To view the power source for a server, select the server on the Servers panel, and then select **Power** from the lower Navigator. The Power Source panel shows the power source configured on the server when HA was configured. You can add or remove power sources on the server, and can power the server on or off, or reset the server.

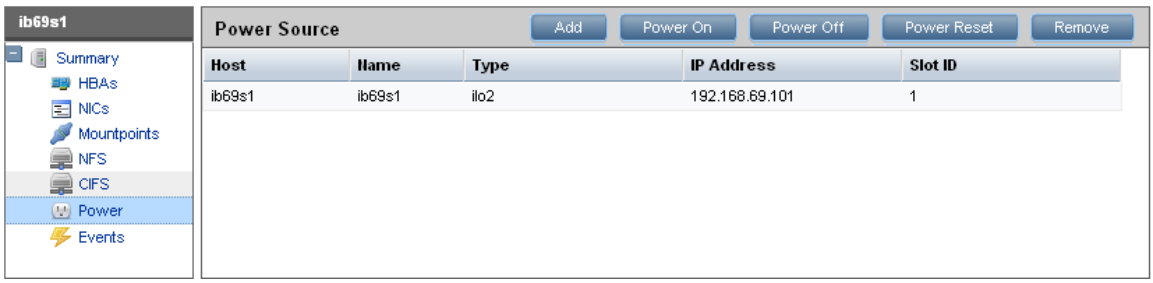

Configuring automated failover manually

To configure automated failover manually, complete these steps:

1. Configure file serving nodes in backup pairs.
2. Identify power sources for the servers in the backup pair.
3. Configure NIC monitoring.
4. Enable automated failover.

1. Configure server backup pairs

File serving nodes are configured in backup pairs, where each server in a pair is the backup for the other. This step is typically done when the cluster is installed. The following restrictions apply:

- The same file system must be mounted on both servers in the pair and the servers must see the same storage.
- In a SAN environment, a server and its backup must use the same storage infrastructure to access a segment's physical volumes (for example, a multiported RAID array).

For a cluster using the unified network configuration, assign backup nodes for the `bond0:1` interface. For example, node1 is the backup for node2, and node2 is the backup for node1.

1. Add the VIF:

   ```
   ibrix_nic -a -n bond0:2 -h node1,node2,node3,node4
   ```

2. Set up a standby server for each VIF:

   ```
   # ibrix_nic -b -H node1/bond0:1,node2/bond0:2
   ibrix_nic -b -H node2/bond0:1,node1/bond0:2
   ibrix_nic -b -H node3/bond0:1,node4/bond0:2
   ibrix_nic -b -H node4/bond0:1,node3/bond0:2
   ```

2. Identify power sources

To implement automated failover, perform a forced manual failover, or remotely power a file serving node up or down, you must set up programmable power sources for the nodes and their backups. Using programmable power sources prevents a "split-brain scenario" between a failing file serving node and its backup, allowing the failing server to be centrally powered down by the Fusion Manager in the case of automated failover, and manually in the case of a forced manual failover.

StoreAll software works with iLO, IPMI, OpenIPMI, and OpenIPMI2 integrated power sources.

The following configuration steps are required when setting up integrated power sources:

- For automated failover, ensure that the Fusion Manager has LAN access to the power sources.
- Install the environment and any drivers and utilities, as specified by the vendor documentation. If you plan to protect access to the power sources, set up the UID and password to be used.

Use the following command to identify a power source:

```
ibrix_powersrc -a -t {ipmi|openipmi|openipmi2|ilo} -h HOSTNAME -I IPADDR
-u USERNAME -p PASSWORD
```

For example, to identify an iLO power source at IP address 192.168.3.170 for node `ss01`:

```
ibrix_powersrc -a -t ilo -h ss01 -I 192.168.3.170 -u Administrator -p
password
```

3. Configure NIC monitoring

NIC monitoring should be configured on user VIFs that will be used by NFS, SMB, FTP, or HTTP.

IMPORTANT: When configuring NIC monitoring, use the same backup pairs that you used when configuring backup servers.

Identify the servers in a backup pair as NIC monitors for each other. Because the monitoring must be declared in both directions, enter a separate command for each server in the pair.

```
ibrix_nic -m -h MONHOST -A DESTHOST/IFNAME
```

The following example sets up monitoring for NICs over `bond0:1`:

```
ibric_nic -m -h node1 -A node2/bond0:1
ibric_nic -m -h node2 -A node1/bond0:1
ibric_nic -m -h node3 -A node4/bond0:1
ibric_nic -m -h node4 -A node3/bond0:1
```

The next example sets up server `s2.hp.com` to monitor server `s1.hp.com` over user network interface `eth1`:

```
ibrix_nic -m -h s2.hp.com -A s1.hp.com/eth1
```

4. Enable automated failover

Automated failover is turned off by default. When automated failover is turned on, the Fusion Manager starts monitoring heartbeat messages from file serving nodes. You can turn automated failover on and off for all file serving nodes or for selected nodes.

Turn on automated failover:

```
ibrix_server -m [-h SERVERNAME]
```

Changing the HA configuration manually

Update a power source:

If you change the IP address or password for a power source, you must update the configuration database with the changes. The user name and password options are needed only for remotely managed power sources. Include the `-s` option to have the Fusion Manager skip BMC.

```
ibrix_powersrc -m [-I IPADDR] [-u USERNAME] [-p PASSWORD]  [-s] -h POWERSRCLIST
```

The following command changes the IP address for power source `ps1`:

```
ibrix_powersrc -m -I 192.168.3.153 -h ps1
```

Disassociate a server from a power source:

You can dissociate a file serving node from a power source by dissociating it from slot 1 (its default association) on the power source. Use the following command:

```
ibrix_hostpower -d -s POWERSOURCE -h HOSTNAME
```

Delete a power source:

To conserve storage, delete power sources that are no longer in use. If you are deleting multiple power sources, use commas to separate them.

```
ibrix_powersrc -d -h POWERSRCLIST
```

Delete NIC monitoring:

To delete NIC monitoring, use the following command:

```
ibrix_nic -m -h MONHOST -D DESTHOST/IFNAME
```

Delete NIC standbys:

To delete a standby for a NIC, use the following command:

```
ibrix_nic -b -U HOSTNAME1/IFNAME1
```

For example, to delete the standby that was assigned to interface `eth2` on file serving node `s1.hp.com`:

```
ibrix_nic -b -U s1.hp.com/eth2
```

Turn off automated failover:

```
ibrix_server -m -U [-h SERVERNAME]
```

To specify a single file serving node, include the `-h SERVERNAME` option.

Failing a server over manually

The server to be failed over must belong to a backup pair. The server can be powered down or remain up during the procedure. You can perform a manual failover at any time, regardless of whether automated failover is in effect. Manual failover does not require the use of a programmable power supply. However, if you have identified a power supply for the server, you can power it down before the failover.

Use the GUI or the CLI to fail over a file serving node:

- On the GUI, select the node on the Servers panel and then click **Failover** on the Summary panel.
- On the CLI, run `ibrix_server -f`, specifying the node to be failed over as the *HOSTNAME*. If appropriate, include the `-p` option to power down the node before segments are migrated:

  ```
  ibrix_server -f [-p] -h HOSTNAME
  ```

Check the Summary panel or run the following command to determine whether the failover was successful:

```
ibrix_server -l
```

The STATE field indicates the status of the failover. If the field persistently shows `Down-InFailover` or `Up-InFailover`, the failover did not complete; contact HP Support for assistance. For information about the values that can appear in the STATE field, see "What happens during a failover" (page 55).

Failing back a server

After an automated or manual failover of a server, you must manually fail back the server, which restores ownership of the failed-over segments and network interfaces to the server. Before failing back the server, confirm that it can see all of its storage resources and networks. The segments owned by the server will not be accessible if the server cannot see its storage.

To fail back a node from the GUI, select the node on the Servers panel and then click **Failback** on the Summary panel.

On the GUI, select the node on the Servers panel and then click **Failback** on the Summary pane

On the CLI, run the following command, where *HOSTNAME* is the failed-over node:

```
ibrix_server -f -U -h HOSTNAME
```

After failing back the node, check the Summary panel or run the `ibrix_server -l` command to determine whether the failback completed fully. If the failback is not complete, contact HP Support.

NOTE: A failback might not succeed if the time period between the failover and the failback is too short, and the primary server has not fully recovered. HP recommends ensuring that both servers are up and running and then waiting 60 seconds before starting the failback. Use the `ibrix_server -l` command to verify that the primary server is up and running. The status should be `Up-FailedOver` before performing the failback.

Setting up HBA monitoring

You can configure High Availability to initiate automated failover upon detection of a failed HBA. HBA monitoring can be set up for either dual-port HBAs with built-in standby switching or single-port HBAs, whether standalone or paired for standby switching via software. The StoreAll software

does not play a role in vendor- or software-mediated HBA failover; traffic moves to the remaining functional port with no Fusion Manager involvement.

HBAs use worldwide names for some parameter values. These are either worldwide node names (WWNN) or worldwide port names (WWPN). The WWPN is the name an HBA presents when logging in to a SAN fabric. Worldwide names consist of 16 hexadecimal digits grouped in pairs. In StoreAll software, these are written as dot-separated pairs (for example, 21.00.00.e0.8b.05.05.04).

To set up HBA monitoring, first discover the HBAs, and then perform the procedure that matches your HBA hardware:

- For single-port HBAs without built-in standby switching: Turn on HBA monitoring for all ports that you want to monitor for failure.
- For dual-port HBAs with built-in standby switching and single-port HBAs that have been set up as standby pairs in a software operation: Identify the standby pairs of ports to the configuration database and then turn on HBA monitoring for all paired ports. If monitoring is turned on for just one port in a standby pair and that port fails, the Fusion Manager will fail over the server even though the HBA has automatically switched traffic to the surviving port. When monitoring is turned on for both ports, the Fusion Manager initiates failover only when both ports in a pair fail.

When both HBA monitoring and automated failover for file serving nodes are configured, the Fusion Manager will fail over a server in two situations:

- **Both ports in a monitored set of standby-paired ports fail**. Because all standby pairs were identified in the configuration database, the Fusion Manager knows that failover is required only when both ports fail.
- **A monitored single-port HBA fails.** Because no standby has been identified for the failed port, the Fusion Manager knows to initiate failover immediately.

Discovering HBAs

You must discover HBAs before you set up HBA monitoring, when you replace an HBA, and when you add a new HBA to the cluster. Discovery adds the WWPN for the port to the configuration database.

```
ibrix_hba -a [-h HOSTLIST]
```

Adding standby-paired HBA ports

Identifying standby-paired HBA ports to the configuration database allows the Fusion Manager to apply the following logic when they fail:

- If one port in a pair fails, do nothing. Traffic will automatically switch to the surviving port, as configured by the HBA vendor or the software.
- If both ports in a pair fail, fail over the server's segments to the standby server.

Use the following command to identify two HBA ports as a standby pair:

```
ibrix_hba -b -P WWPN1:WWPN2 -h HOSTNAME
```

Enter the WWPN as decimal-delimited pairs of hexadecimal digits. The following command identifies port 20.00.12.34.56.78.9a.bc as the standby for port 42.00.12.34.56.78.9a.bc for the HBA on file serving node s1.hp.com:

```
ibrix_hba -b -P 20.00.12.34.56.78.9a.bc:42.00.12.34.56.78.9a.bc -h s1.hp.com
```

Turning HBA monitoring on or off

If your cluster uses single-port HBAs, turn on monitoring for all of the ports to set up automated failover in the event of HBA failure. Use the following command:

```
ibrix_hba -m -h HOSTNAME -p PORT
```

For example, to turn on HBA monitoring for port 20.00.12.34.56.78.9a.bc on node s1.hp.com:

```
ibrix_hba -m -h s1.hp.com -p 20.00.12.34.56.78.9a.bc
```

To turn off HBA monitoring for an HBA port, include the `-U` option:

```
ibrix_hba -m -U -h HOSTNAME -p PORT
```

Deleting standby port pairings

Deleting port pairing information from the configuration database does not remove the standby pairing of the ports. The standby pairing is either built in by the HBA vendor or implemented by software.

To delete standby-paired HBA ports from the configuration database, enter the following command:

```
ibrix_hba -b -U -P WWPN1:WWPN2 -h HOSTNAME
```

For example, to delete the pairing of ports 20.00.12.34.56.78.9a.bc and 42.00.12.34.56.78.9a.bc on node s1.hp.com:

```
ibrix_hba -b -U -P 20.00.12.34.56.78.9a.bc:42.00.12.34.56.78.9a.bc
-h s1.hp.com
```

Deleting HBAs from the configuration database

Before switching an HBA to a different machine, delete the HBA from the configuration database:

```
ibrix_hba -d -h HOSTNAME -w WWNN
```

Displaying HBA information

Use the following command to view information about the HBAs in the cluster. To view information for all hosts, omit the `-h HOSTLIST` argument.

```
ibrix_hba -l [-h HOSTLIST]
```

The output includes the following fields:

Field	Description
Host	Server on which the HBA is installed.
Node WWN	This HBA's WWNN.
Port WWN	This HBA's WWPN.
Port State	Operational state of the port.
Backup Port WWN	WWPN of the standby port for this port (standby-paired HBAs only).
Monitoring	Whether HBA monitoring is enabled for this port.

Checking the High Availability configuration

Use the `ibrix_haconfig` command to determine whether High Availability features have been configured for specific file serving nodes. The command checks for the following features and provides either a summary or a detailed report of the results:

- Programmable power source
- Standby server or standby segments
- Cluster and user network interface monitors
- Standby network interface for each user network interface
- HBA port monitoring
- Status of automated failover (on or off)

For each High Availability feature, the summary report returns status for each tested file serving node and optionally for their standbys:

- **Passed.** The feature has been configured.
- **Warning.** The feature has not been configured, but the significance of the finding is not clear. For example, the absence of discovered HBAs can indicate either that the HBA monitoring feature was not configured or that HBAs are not physically present on the tested servers.
- **Failed.** The feature has not been configured.

The detailed report includes an overall result status for all tested file serving nodes and describes details about the checks performed on each High Availability feature. By default, the report includes details only about checks that received a Failed or a Warning result. You can expand the report to include details about checks that received a Passed result.

Viewing a summary report

Use the `ibrix_haconfig -l` command to see a summary of all file serving nodes. To check specific file serving nodes, include the `-h` *`HOSTLIST`* argument. To check standbys, include the `-b` argument. To view results only for file serving nodes that failed a check, include the `-f` argument.

```
ibrix_haconfig -l [-h HOSTLIST] [-f] [-b]
```

For example, to view a summary report for file serving nodes xs01.hp.com and xs02.hp.com:

```
ibrix_haconfig -l -h xs01.hp.com,xs02.hp.com

Host             HA Configuration Power Sources Backup Servers Auto Failover
  Nics Monitored  Standby Nics HBAs Monitored
xs01.hp.com      FAILED           PASSED        PASSED          PASSED
  FAILED          PASSED       FAILED
xs02.hp.com      FAILED           PASSED        FAILED          FAILED
  FAILED          WARNED       WARNED
```

Viewing a detailed report

Execute the `ibrix_haconfig -i` command to view the detailed report:

```
ibrix_haconfig -i [-h HOSTLIST] [-f] [-b] [-s] [-v]
```

The `-h` *`HOSTLIST`* option lists the nodes to check. To also check standbys, include the `-b` option. To view results only for file serving nodes that failed a check, include the `-f` argument. The `-s` option expands the report to include information about the file system and its segments. The `-v` option produces detailed information about configuration checks that received a Passed result.

For example, to view a detailed report for file serving node xs01.hp.com:

```
ibrix_haconfig -i -h xs01.hp.com

--------------- Overall HA Configuration Checker Results ---------------
FAILED
--------------- Overall Host Results ---------------
Host        HA Configuration Power Sources Backup Servers Auto Failover Nics Monitored
        Standby Nics  HBAs Monitored
xs01.hp.com FAILED           PASSED        PASSED         PASSED        FAILED
        PASSED        FAILED

--------------- Server xs01.hp.com FAILED Report ---------------

Check Description                                 Result  Result Information
================================================= ======  ==================
Power source(s) configured                        PASSED
Backup server or backups for segments configured PASSED
Automatic server failover configured              PASSED

Cluster & User Nics monitored
  Cluster nic xs01.hp.com/eth1 monitored          FAILED   Not monitored

User nics configured with a standby nic           PASSED

HBA ports monitored
```

```
Hba port 21.01.00.e0.8b.2a.0d.6d monitored     FAILED   Not monitored
Hba port 21.00.00.e0.8b.0a.0d.6d monitored     FAILED   Not monitored
```

Capturing a core dump from a failed node

The crash capture feature collects a core dump from a failed node when the Fusion Manager initiates failover of the node. You can use the core dump to analyze the root cause of the node failure. When enabled, crash capture is supported for both automated and manual failover. Failback is not affected by this feature. By default, crash capture is disabled. This section provides the prerequisites and steps for enabling crash capture.

NOTE: Enabling crash capture adds a delay (up to 240 seconds) to the failover to allow the crash kernel to load. The failover process ensures that the crash kernel is loaded before continuing.

When crash capture is enabled, the system takes the following actions when a node fails:

1. The Fusion Manager triggers a core dump on the failed node when failover starts, changing the state of the node to `Up, InFailover`.
2. The failed node boots into the crash kernel. The state of the node changes to `Dumping, InFailover`.
3. The failed node continues with the failover, changing state to `Dumping, FailedOver`.
4. After the core dump is created, the failed node reboots and its state changes to `Up, FailedOver`.

IMPORTANT: Complete the steps in "Prerequisites for setting up the crash capture" (page 68) before setting up the crash capture.

Prerequisites for setting up the crash capture

The following parameters must be configured in the ROM-based setup utility (RBSU) before a crash can be captured automatically on a file server node in failed condition.

1. Start RBSU – Reboot the server, and then Press **F9** Key.
2. Highlight the System Options option in main menu, and then press the **Enter** key. Highlight the Virtual Serial Port option (below figure), and then press the **Enter** key. Select the **COM1** port, and then press the **Enter** key.

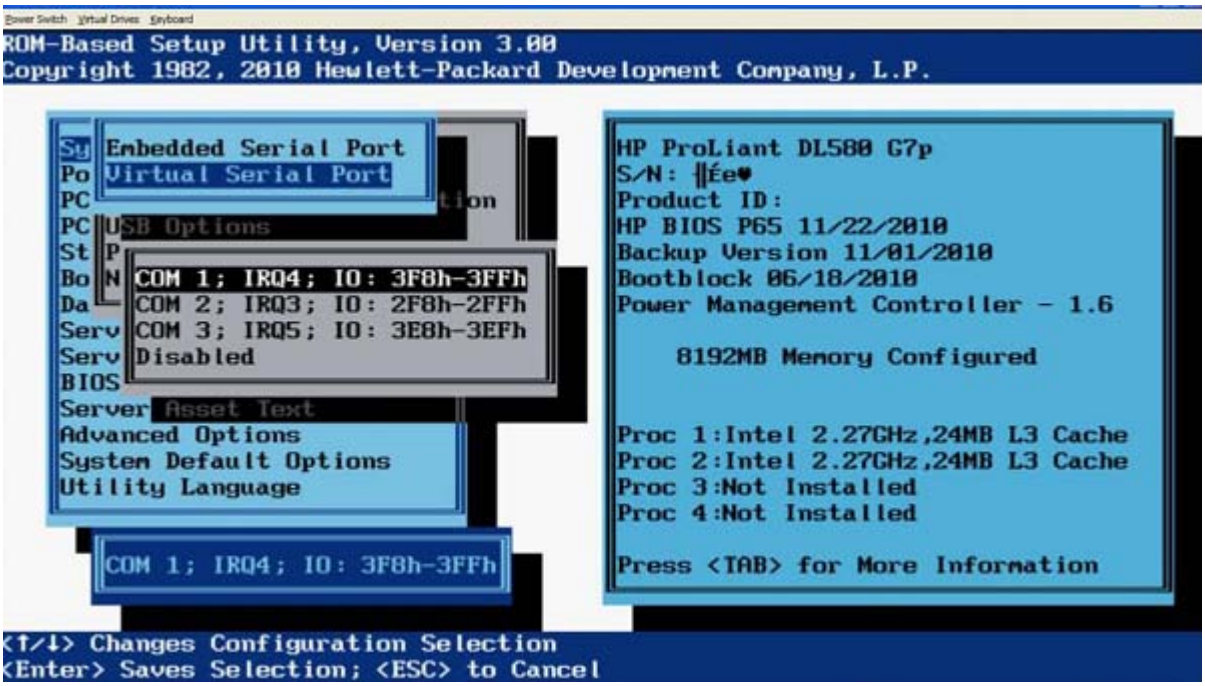

3. Highlight the BIOS Serial Console & EMS option in main menu, and then press the Enter key. Highlight the BIOS Serial Console Port option and then press the Enter key. Select the **COM1** port, and then press the Enter key.
4. Highlight the **BIOS Serial Console Baud Rate** option, and then press the Enter key. Select the **115200** Serial Baud Rate.
5. Highlight the Server Availability option in main menu, and then press the Enter key. Highlight the ASR Timeout option and then press the Enter key. Select the **30 Minutes**, and then press the Enter key.
6. To exit RBSU, press **Esc** until the main menu is displayed. Then, at the main menu, press **F10**. The server automatically restarts.

Setting up nodes for crash capture

IMPORTANT: Complete the steps in "Prerequisites for setting up the crash capture" (page 68) before starting the steps in this section.

To set up nodes for crash capture, complete the following steps:

1. Enable crash capture. Run the following command:

   ```
   ibrix_host_tune -S { -h HOSTLIST | -g GROUPLIST } -o
   trigger_crash_on_failover=1
   ```

2. Tune Fusion Manager to set the DUMPING status timeout by entering the following command:

   ```
   ibrix_fm_tune -S -o dumpingStatusTimeout=240
   ```

   This command is required to delay the failover until the crash kernel is loaded; otherwise, Fusion Manager will bring down the failed node.

6 Configuring cluster event notification

Cluster events

There are three categories for cluster events:

	Alerts. Disruptive events that can result in loss of access to file system data.
	Warnings. Potentially disruptive conditions where file system access is not lost, but if the situation is not addressed, it can escalate to an alert condition.
	Information. Normal events that change the cluster.

The following table lists examples of events included in each category.

Event Type	Trigger Point	Name
ALERT	User fails to log into GUI	`login.failure`
	File system is unmounted	`filesystem.unmounted`
	File serving node is down/restarted	`server.status.down`
	File serving node terminated unexpectedly	`server.unreachable`
WARN	User migrates segment using GUI	`segment.migrated`
INFO	User successfully logs in to GUI	`login.success`
	File system is created	`filesystem.cmd`
	File serving node is deleted	`server.deregistered`
	NIC is added using GUI	`nic.added`
	NIC is removed using GUI	`nic.removed`
	Physical storage is discovered and added using management console	`physicalvolume.added`
	Physical storage is deleted using management console	`physicalvolume.deleted`

You can be notified of cluster events by email or SNMP traps. To view the list of supported events, use the command `ibrix_event -q`.

NOTE: The StoreAll event system does not report events from the MSA array. Instead, configure event notification using the SMU on the array. For more information, see "Event notification for MSA array systems" (page 75).

Setting up email notification of cluster events

You can set up event notifications by event type or for one or more specific events. To set up automatic email notification of cluster events, associate the events with email recipients and then configure email settings to initiate the notification process.

Associating events and email addresses

You can associate any combination of cluster events with email addresses: all Alert, Warning, or Info events, all events of one type plus a subset of another type, or a subset of all types.

The notification threshold for Alert events is 90% of capacity. Threshold-triggered notifications are sent when a monitored system resource exceeds the threshold and are reset when the resource utilization dips 10% below the threshold. For example, a notification is sent the first time usage reaches 90% or more. The next notice is sent only if the usage declines to 80% or less (event is reset), and subsequently rises again to 90% or above.

To associate all types of events with recipients, omit the `-e` argument in the following command:

```
ibrix_event -c [-e ALERT|WARN|INFO|EVENTLIST] -m EMAILLIST
```

Use the `ALERT`, `WARN`, and `INFO` keywords to make specific type associations or use *`EVENTLIST`* to associate specific events.

The following command associates all types of events to admin@hp.com:

```
ibrix_event -c -m admin@hp.com
```

The next command associates all Alert events and two Info events to admin@hp.com:

```
ibrix_event -c -e ALERT,server.registered,filesystem.space.full
-m admin@hp.com
```

Configuring email notification settings

To configure email notification settings, specify the SMTP server and header information and turn the notification process on or off.

```
ibrix_event -m on|off -s SMTP -f from [-r reply-to] [-t subject]
```

The server must be able to receive and send email and must recognize the From and Reply-to addresses. Be sure to specify valid email addresses, especially for the SMTP server. If an address is not valid, the SMTP server will reject the email.

The following command configures email settings to use the mail.hp.com SMTP server and turns on notifications:

```
ibrix_event -m on -s mail.hp.com -f FM@hp.com -r MIS@hp.com -t Cluster1 Notification
```

NOTE: The state of the email notification process has no effect on the display of cluster events in the GUI.

Dissociating events and email addresses

To remove the association between events and email addresses, use the following command:

```
ibrix_event -d [-e ALERT|WARN|INFO|EVENTLIST] -m EMAILLIST
```

For example, to dissociate event notifications for admin@hp.com:

```
ibrix_event -d -m admin@hp.com
```

To turn off all Alert notifications for admin@hp.com:

```
ibrix_event -d -e ALERT -m admin@hp.com
```

To turn off the `server.registered` and `filesystem.created` notifications for admin1@hp.com and admin2@hp.com:

```
ibrix_event -d -e server.registered,filesystem.created -m admin1@hp.com,admin2@hp.com
```

Testing email addresses

To test an email address with a test message, notifications must be turned on. If the address is valid, the command signals success and sends an email containing the settings to the recipient. If the address is not valid, the command returns an `address failed` exception.

```
ibrix_event -u -n EMAILADDRESS
```

Viewing email notification settings

The `ibrix_event -L` command provides comprehensive information about email settings and configured notifications.

```
ibrix_event -L
Email Notification  :  Enabled
SMTP Server         :  mail.hp.com
From                :  FM@hp.com
Reply To            :  MIS@hp.com

EVENT                                  LEVEL  TYPE   DESTINATION
-------------------------------------  -----  -----  -----------
asyncrep.completed                     ALERT  EMAIL  admin@hp.com
asyncrep.failed                        ALERT  EMAIL  admin@hp.com
```

Setting up SNMP notifications

The StoreAll software supports SNMP (Simple Network Management Protocol) V1, V2, and V3.

Whereas SNMPV2 security was enforced by use of community password strings, V3 introduces the USM and VACM. Discussion of these models is beyond the scope of this document. Refer to RFCs 3414 and 3415 at http://www.ietf.org for more information. Note the following:

- In the SNMPV3 environment, every message contains a user name. The function of the USM is to authenticate users and ensure message privacy through message encryption and decryption. Both authentication and privacy, and their passwords, are optional and will use default settings where security is less of a concern.
- With users validated, the VACM determines which managed objects these users are allowed to access. The VACM includes an access scheme to control user access to managed objects; context matching to define which objects can be accessed; and MIB views, defined by subsets of IOD subtree and associated bitmask entries, which define what a particular user can access in the MIB.

Steps for setting up SNMP include:

- Agent configuration (all SNMP versions)
- Trapsink configuration (all SNMP versions)
- Associating event notifications with trapsinks (all SNMP versions)
- View definition (V3 only)
- Group and user configuration (V3 only)

StoreAll software implements an SNMP agent that supports the private StoreAll software MIB. The agent can be polled and can send SNMP traps to configured trapsinks.

Setting up SNMP notifications is similar to setting up email notifications. You must associate events to trapsinks and configure SNMP settings for each trapsink to enable the agent to send a trap when an event occurs.

NOTE: When Phone Home is enabled, you cannot edit or change the configuration of the StoreAll SNMP agent with the `ibrix_snmpagent`. However, you can add trapsink IPs with `ibrix_snmptrap` and can associate events to the trapsink IP with `ibrix_event`.

Configuring the SNMP agent

The SNMP agent is created automatically when the Fusion Manager is installed. It is initially configured as an SNMPv2 agent and is off by default.

Some SNMP parameters and the SNMP default port are the same, regardless of SNMP version. The default agent port is 161. *SYSCONTACT*, *SYSNAME*, and *SYSLOCATION* are optional MIB-II agent parameters that have no default values.

NOTE: The default SNMP agent port was changed from 5061 to 161 in the StoreAll 6.1 release. This port number cannot be changed.

The `-c` and `-s` options are also common to all SNMP versions. The `-c` option turns the encryption of community names and passwords on or off. There is no encryption by default. Using the `-s` option toggles the agent on and off; it turns the agent on by starting a listener on the SNMP port, and turns it off by shutting off the listener. The default is off.

The format for a v1 or v2 update command follows:

```
ibrix_snmpagent -u -v {1|2} [-p PORT] [-r READCOMMUNITY] [-w WRITECOMMUNITY]
[-t SYSCONTACT] [-n SYSNAME] [-o SYSLOCATION] [-c {yes|no}] [-s {on|off}]
```

The update command for SNMPv1 and v2 uses optional community names. By convention, the default *READCOMMUNITY* name used for read-only access and assigned to the agent is `public`. No default *WRITECOMMUNITY* name is set for read-write access (although the name `private` is often used).

The following command updates a v2 agent with the write community name `private`, the agent's system name, and that system's physical location:

```
ibrix_snmpagent -u -v 2 -w private -n agenthost.domain.com -o DevLab-B3-U6
```

The SNMPv3 format adds an optional *engine id* that overrides the default value of the agent's host name. The format also provides the `-y` and `-z` options, which determine whether a v3 agent can process v1/v2 read and write requests from the management station. The format is:

```
ibrix_snmpagent -u -v 3 [-e engineId] [-p PORT] [-r READCOMMUNITY]
[-w WRITECOMMUNITY] [-t SYSCONTACT] [-n SYSNAME] [-o SYSLOCATION]
[-y {yes|no}] [-z {yes|no}] [-c {yes|no}] [-s {on|off}]
```

Configuring trapsink settings

A *trapsink* is the host destination where agents send *traps*, which are asynchronous notifications sent by the agent to the management station. A trapsink is specified either by name or IP address. StoreAll software supports multiple trapsinks; you can define any number of trapsinks of any SNMP version, but you can define only one trapsink per host, regardless of the version.

At a minimum, trapsink configuration requires a destination host and SNMP version. All other parameters are optional and many assume the default value if no value is specified.

The format for creating a v1/v2 trapsink is:

```
ibrix_snmptrap -c -h HOSTNAME -v {1|2} [-p PORT] [-m COMMUNITY] [-s {on|off}]
```

If a port is not specified, the command defaults to port 162. If a community is not specified, the command defaults to the community name `public`. The `-s` option toggles agent trap transmission on and off. The default is on. For example, to create a v2 trapsink with a new community name, enter:

```
ibrix_snmptrap -c -h lab13-116 -v 2 -m private
```

For a v3 trapsink, additional options define security settings. *USERNAME* is a v3 user defined on the trapsink host and is required. The security level associated with the trap message depends on which passwords are specified—the authentication password, both the authentication and privacy passwords, or no passwords. The *CONTEXT_NAME* is required if the trap receiver has defined subsets of managed objects. The format is:

```
ibrix_snmptrap -c -h HOSTNAME -v 3 [-p PORT] -n USERNAME [-j {MD5|SHA}]
[-k AUTHORIZATION_PASSWORD] [-y {DES|AES}] [-z PRIVACY_PASSWORD]
[-x CONTEXT_NAME] [-s {on|off}]
```

The following command creates a v3 trapsink with a named user and specifies the passwords to be applied to the default algorithms. If specified, passwords must contain at least eight characters.

```
ibrix_snmptrap -c -h lab13-114 -v 3 -n trapsender -k auth-passwd -z priv-passwd
```

Associating events and trapsinks

Associating events with trapsinks is similar to associating events with email recipients, except that you specify the host name or IP address of the trapsink instead of an email address.

Use the `ibrix_event` command to associate SNMP events with trapsinks. The format is:

```
ibrix_event -c -y SNMP [-e ALERT|INFO|EVENTLIST]
 -m TRAPSINK
```

For example, to associate all Alert events and two Info events with a trapsink at IP address 192.168.2.32, enter:

```
ibrix_event -c -y SNMP -e ALERT,server.registered,
filesystem.created -m 192.168.2.32
```

Use the `ibrix_event -d` command to dissociate events and trapsinks:

```
ibrix_event -d -y SNMP [-e ALERT|INFO|EVENTLIST] -m TRAPSINK
```

Defining views

A MIB view is a collection of paired OID subtrees and associated bitmasks that identify which subidentifiers are significant to the view's definition. Using the bitmasks, individual OID subtrees can be included in or excluded from the view.

An instance of a managed object belongs to a view if:

- The OID of the instance has at least as many sub-identifiers as the OID subtree in the view.
- Each sub-identifier in the instance and the subtree match when the bitmask of the corresponding sub-identifier is nonzero.

The Fusion Manager automatically creates the `excludeAll` view that blocks access to all OIDs. This view cannot be deleted; it is the default read and write view if one is not specified for a group with the `ibrix_snmpgroup` command. The catch-all OID and mask are:

```
OID = .1
Mask = .1
```

Consider these examples, where instance .1.3.6.1.2.1.1 matches, instance .1.3.6.1.4.1 matches, and instance .1.2.6.1.2.1 does not match.

```
OID = .1.3.6.1.4.1.18997
Mask = .1.1.1.1.1.1.1

OID = .1.3.6.1.2.1
Mask = .1.1.0.1.0.1
```

To add a pairing of an OID subtree value and a mask value to a new or existing view, use the following format:

```
ibrix_snmpview -a -v VIEWNAME [-t {include|exclude}] -o OID_SUBTREE [-m MASK_BITS]
```

The subtree is added in the named view. For example, to add the StoreAll software private MIB to the view named `hp`, enter:

```
ibrix_snmpview -a -v hp -o .1.3.6.1.4.1.18997 -m .1.1.1.1.1.1.1
```

Configuring groups and users

A *group* defines the access control policy on managed objects for one or more users. All users must belong to a group. Groups and users exist only in SNMPv3. Groups are assigned a security level, which enforces use of authentication and privacy, and specific read and write views to identify which managed objects group members can read and write.

The command to create a group assigns its SNMPv3 security level, read and write views, and context name. A *context* is a collection of managed objects that can be accessed by an SNMP entity. A related option, `-m`, determines how the context is matched. The format follows:

```
ibrix_snmpgroup -c -g GROUPNAME [-s {noAuthNoPriv|authNoPriv|authPriv}]
[-r READVIEW] [-w WRITEVIEW]
```

For example, to create the group `group2` to require authorization, no encryption, and read access to the `hp` view, enter:

```
ibrix_snmpgroup -c -g group2 -s authNoPriv -r hp
```

The format to create a user and add that user to a group follows:

```
ibrix_snmpuser -c -n USERNAME -g GROUPNAME [-j {MD5|SHA}]
[-k AUTHORIZATION_PASSWORD] [-y {DES|AES}] [-z PRIVACY_PASSWORD]
```

Authentication and privacy settings are optional. An authentication password is required if the group has a security level of either authNoPriv or authPriv. The privacy password is required if the group has a security level of authPriv. If unspecified, MD5 is used as the authentication algorithm and DES as the privacy algorithm, with no passwords assigned.

For example, to create `user3`, add that user to `group2`, and specify an authorization password for authorization and no encryption, enter:

```
ibrix_snmpuser -c -n user3 -g group2 -k auth-passwd -s authNoPriv
```

Deleting elements of the SNMP configuration

All SNMP commands use the same syntax for delete operations, using `-d` to indicate the object is to delete. The following command deletes a list of hosts that were trapsinks:

```
ibrix_snmptrap -d -h lab15-12.domain.com,lab15-13.domain.com,lab15-14.domain.com
```

There are two restrictions on SNMP object deletions:

- A view cannot be deleted if it is referenced by a group.
- A group cannot be deleted if it is referenced by a user.

Listing SNMP configuration information

All SNMP commands employ the same syntax for list operations, using the `-l` flag. For example:

```
ibrix_snmpgroup -l
```

This command lists the defined group settings for all SNMP groups. Specifying an optional group name lists the defined settings for that group only.

Event notification for MSA array systems

The StoreAll event system does not report events for MSA array systems. Instead, configure event notification for the MSA using the SMU configuration wizard. In the SMU Configuration View panel, right-click the system and select either **Configuration** > **Configuration Wizard** or **Wizards** > **Configuration Wizard.** Configure up to four email addresses and three SNMP trap hosts to receive notifications of system events.

In the Email Configuration section, set the options:

- Notification Level. Select the minimum severity for which the system should send notifications: Critical (only); Error (and Critical); Warning (and Error and Critical); Informational (all). The default is `none`, which disables email notification.
- SMTP Server address. The IP address of the SMTP mail server to use for the email messages. If the mail server is not on the local network, make sure that the gateway IP address was set in the network configuration step.
- Sender Name. The sender name that is joined with an @ symbol to the domain name to form the "from" address for remote notification. This name provides a way to identify the system that is sending the notification. The sender name can have a maximum of 31 bytes. Because this name is used as part of an email address, do not include spaces. For example: `Storage-1`. If no sender name is set, a default name is created.

- Sender Domain. The domain name that is joined with an @ symbol to the sender name to form the "from" address for remote notification. The domain name can have a maximum of 31 bytes. Because this name is used as part of an email address, do not include spaces. For example: `MyDomain.com`. If the domain name is not valid, some email servers will not process the mail.
- Email Address fields. Up to four email addresses that the system should send notifications to. Email addresses must use the format `user-name@domain-name`. Each email address can have a maximum of 79 bytes. For example: `Admin@MyDomain.com`.

In the SNMP Configuration section, set the options:

- Notification Level. Select the minimum severity for which the system should send notifications: Critical (only); Error (and Critical); Warning (and Error and Critical); Informational (all). The default is `none`, which disables SNMP notification.
- Read Community. The SNMP read password for your network. This password is also included in traps that are sent. The value is case sensitive; can include letters, numbers, hyphens, and underscores; and can have a maximum of 31 bytes. The default is `public`.
- Write Community. The SNMP write password for your network. The value is case sensitive; can include letters, numbers, hyphens, and underscores; and can have a maximum of 31 bytes. The default is `private`.
- Trap Host Address fields. The IP addresses of up to three host systems that are configured to receive SNMP traps.

See the MSA array documentation for additional information. For HP P2000 G3 MSA systems, see the *HP P2000 G3 MSA System SMU Reference Guide*. For P2000 G2 MSA systems, see the *HP 2000 G2 Modular Smart Array Reference Guide*. To locate these documents, go to http://www.hp.com/support/manuals. On the Manuals page, select **storage** >**Disk Storage Systems** > **P2000/MSA Disk Arrays** >**HP 2000sa G2 Modular Smart Array** or **HP P2000 G3 MSA Array Systems**.

7 Configuring system backups

Backing up the Fusion Manager configuration

The Fusion Manager configuration is automatically backed up whenever the cluster configuration changes. The backup occurs on the node hosting the active Fusion Manager. The backup file is stored at `<ibrixhome>/tmp/fmbackup.zip` on that node.

The active Fusion Manager notifies the passive Fusion Manager when a new backup file is available. The passive Fusion Manager then copies the file to `<ibrixhome>/tmp/fmbackup.zip` on the node on which it is hosted. If a Fusion Manager is in maintenance mode, it will also be notified when a new backup file is created, and will retrieve it from the active Fusion Manager.

You can create an additional copy of the backup file at any time. Run the following command, which creates a `fmbackup.zip` file in the `$IBRIXHOME/log` directory:

```
$IBRIXHOME/bin/db_backup.sh
```

Once each day, a `cron` job rotates the `$IBRIXHOME/log` directory into the `$IBRIXHOME/log/daily` subdirectory. The `cron` job also creates a new backup of the Fusion Manager configuration in both `$IBRIXHOME/tmp` and `$IBRIXHOME/log`.

To force a backup, use the following command:

```
ibrix_fm -B
```

IMPORTANT: You will need the backup file to recover from server failures or to undo unwanted configuration changes. Whenever the cluster configuration changes, be sure to save a copy of `fmbackup.zip` in a safe, remote location such as a node on another cluster.

Using NDMP backup applications

The NDMP backup feature can be used to back up and recover entire StoreAll software file systems or portions of a file system. You can use any supported NDMP backup application to perform the backup and recovery operations. (In NDMP terminology, the backup application is referred to as a Data Management Application, or DMA.) The DMA is run on a management station separate from the cluster and communicates with the cluster's file serving nodes over a configurable socket port.

The NDMP backup feature supports the following:

- NDMP protocol versions 3 and 4
- Two-way NDMP operations
- Three-way NDMP operations between two network storage systems

Each file serving node functions as an NDMP Server and runs the NDMP Server daemon (ndmpd) process. When you start a backup or restore operation on the DMA, you can specify the node and tape device to be used for the operation.

Following are considerations for configuring and using the NDMP feature:

- When configuring your system for NDMP operations, attach your tape devices to a SAN and then verify that the file serving nodes to be used for backup/restore operations can see the appropriate devices.
- When performing backup operations, take snapshots of your file systems and then back up the snapshots.
- When directory tree quotas are enabled, an NDMP restore to the original location fails if the hard quota limit is exceeded. The NDMP restore operation first creates a temporary file and then restores a file to the temporary file. After this succeeds, the restore operation overwrites the existing file (if it present in same destination directory) with the temporary file. When the

hard quota limit for the directory tree has been exceeded, NDMP cannot create a temporary file and the restore operation fails.

Configuring NDMP parameters on the cluster

Certain NDMP parameters must be configured to enable communications between the DMA and the NDMP Servers in the cluster. To configure the parameters on the GUI, select **Cluster Configuration** from the Navigator, and then select **NDMP Backup**. The NDMP Configuration Summary shows the default values for the parameters. Click **Modify** to configure the parameters for your cluster on the Configure NDMP dialog box. See the online help for a description of each field.

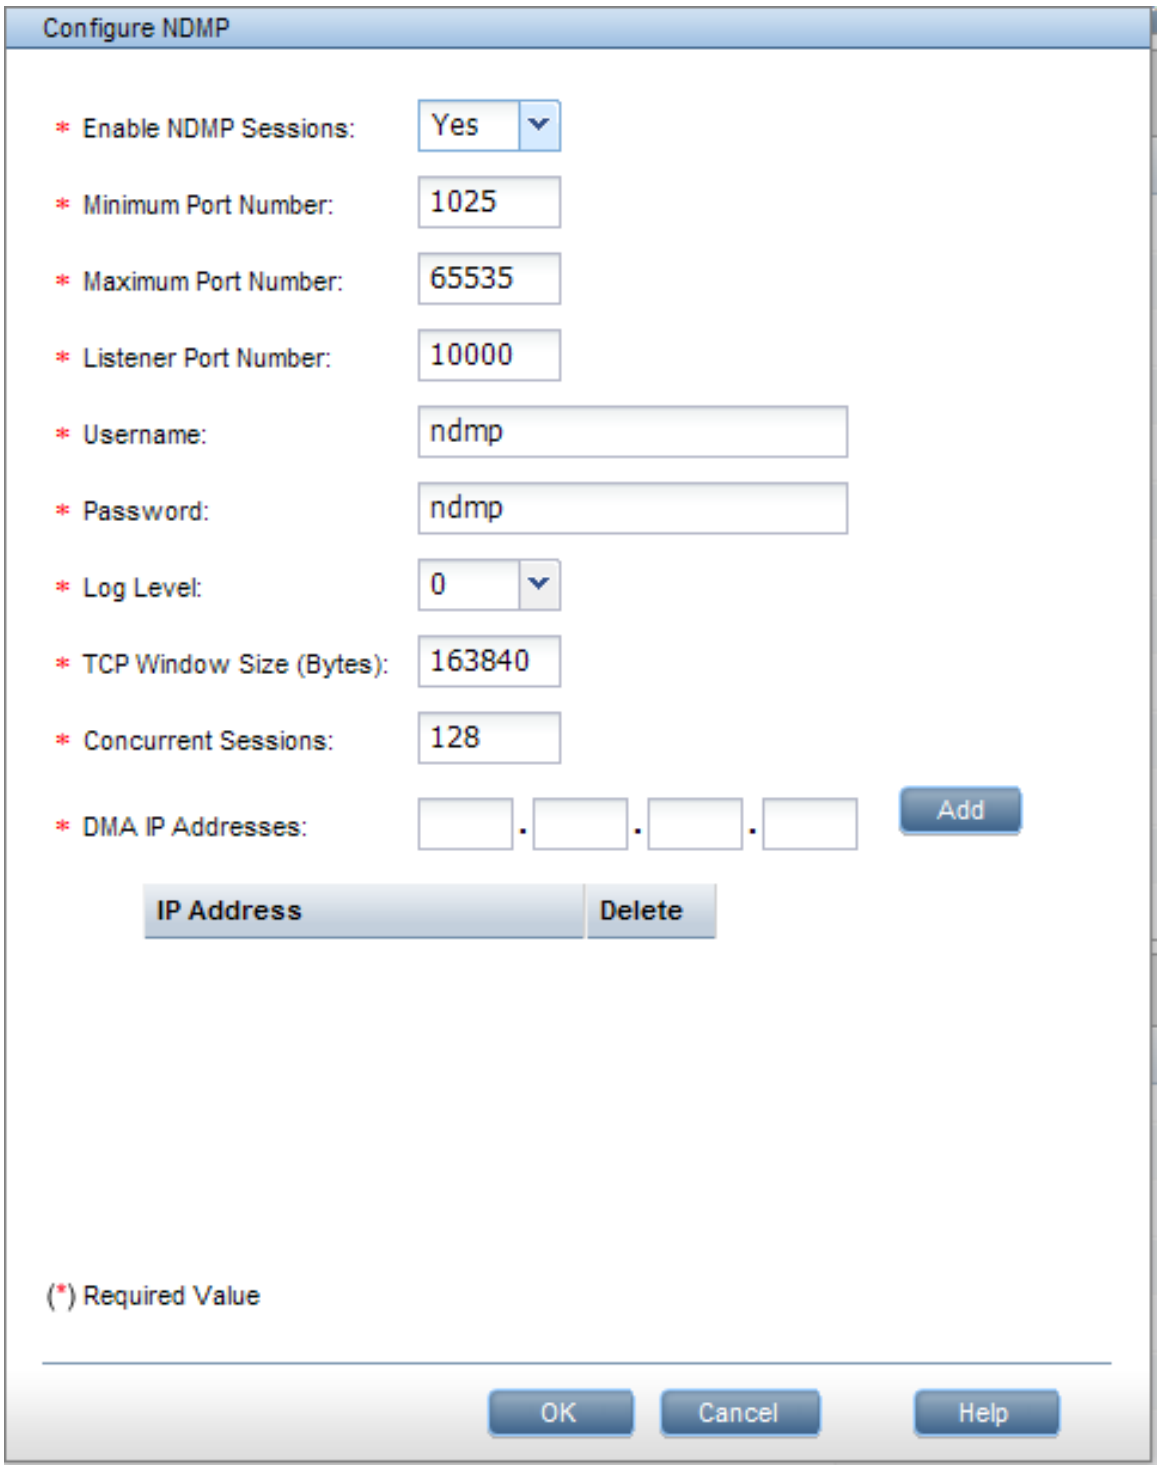

To configure NDMP parameters from the CLI, use the following command:

```
ibrix_ndmpconfig -c [-d IP1,IP2,IP3,...] [-m MINPORT] [-x MAXPORT] [-n LISTENPORT]
[-u USERNAME] [-p PASSWORD] [-e {0=disable,1=enable}] -v [{0=10}] [-w BYTES]
[-z NUMSESSIONS]
```

NDMP process management

All NDMP actions are usually controlled from the DMA. However, if the DMA cannot resolve a problem or you suspect that the DMA may have incorrect information about the NDMP environment, take the following actions from the GUI or CLI:

- Cancel one or more NDMP sessions on a file serving node. Canceling a session stops all spawned sessions processes and frees their resources if necessary.
- Reset the NDMP server on one or more file serving nodes. This step stops all spawned session processes, stops the ndmpd and session monitor daemons, frees all resources held by NDMP, and restarts the daemons.

Viewing or canceling NDMP sessions

To view information about active NDMP sessions, select **Cluster Configuration** from the Navigator, and then select **NDMP Backup** > **Active Sessions**. For each session, the Active NDMP Sessions panel lists the host used for the session, the identifier generated by the backup application, the status of the session (backing up data, restoring data, or idle), the start time, and the IP address used by the DMA.

To cancel a session, select that session and click **Cancel Session**. Canceling a session kills all spawned sessions processes and frees their resources if necessary.

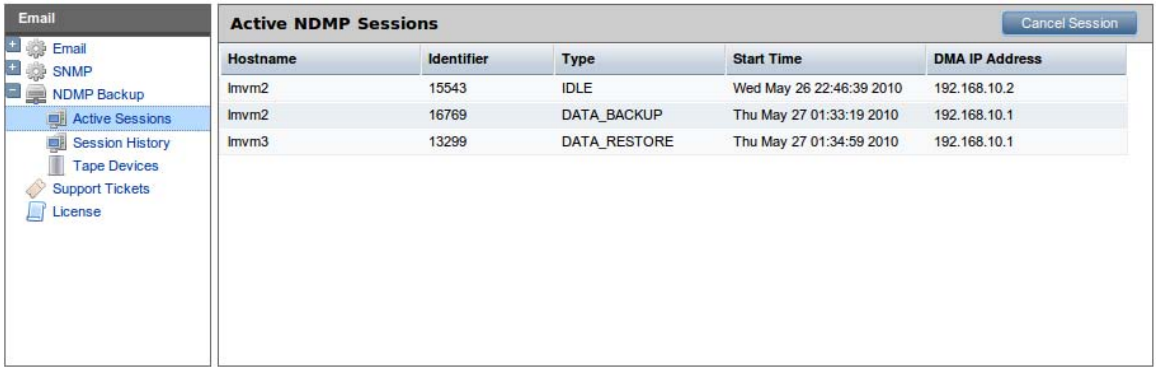

To see similar information for completed sessions, select **NDMP Backup** > **Session History**.

View active sessions from the CLI:

```
ibrix_ndmpsession -l
```

View completed sessions:

```
ibrix_ndmpsession -l -s [-t YYYY-MM-DD]
```

The `-t` option restricts the history to sessions occurring on or before the specified date.

Cancel sessions on a specific file serving node:

```
ibrix_ndmpsession -c SESSION1,SESSION2,SESSION3,... -h HOST
```

Starting, stopping, or restarting an NDMP Server

When a file serving node is booted, the NDMP Server is started automatically. If necessary, you can use the following command to start, stop, or restart the NDMP Server on one or more file serving nodes:

```
ibrix_server -s -t ndmp -c { start | stop | restart} [-h SERVERNAMES]
```

Viewing or rescanning tape and media changer devices

To view the tape and media changer devices currently configured for backups, select **Cluster Configuration** from the Navigator, and then select **NDMP Backup** > **Tape Devices**.

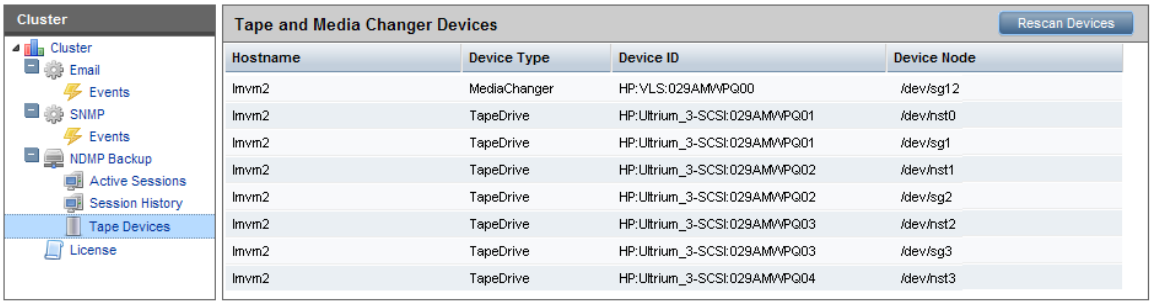

If you add a tape or media changer device to the SAN, click **Rescan Device** to update the list. If you remove a device and want to delete it from the list, reboot all of the servers to which the device is attached.

To view tape and media changer devices from the CLI, use the following command:

```
ibrix_tape -l
```

To rescan for devices, use the following command:

```
ibrix_tape -r
```

NDMP events

An NDMP Server can generate three types of events: INFO, WARN, and ALERT. These events are displayed on the GUI and can be viewed with the `ibrix_event` command.

INFO events. Identifies when major NDMP operations start and finish, and also report progress. For example:

```
7012:Level 3 backup of /mnt/ibfs7 finished at Sat Nov  7 21:20:58 PST 2011
7013:Total Bytes = 38274665923, Average throughput = 236600391 bytes/sec.
```

WARN events. Indicates an issue with NDMP access, the environment, or NDMP operations. Be sure to review these events and take any necessary corrective actions. Following are some examples:

```
0000:Unauthorized NDMP Client 16.39.40.201 trying to connect
4002:User [joe] md5 mode login failed.
```

ALERT events. Indicates that an NDMP action has failed. For example:

```
1102: Cannot start the session_monitor daemon, ndmpd exiting.
7009:Level 6 backup of /mnt/shares/accounts1 failed (writing eod header error).
8001:Restore Failed to read data stream signature.
```

You can configure the system to send email or SNMP notifications when these types of events occur.

8 Creating host groups for StoreAll clients

A *host group* is a named set of StoreAll clients. Host groups provide a convenient way to centrally manage clients. You can put different sets of clients into host groups and then perform the following operations on all members of the group:

- Create and delete mount points
- Mount file systems
- Prefer a network interface
- Tune host parameters
- Set allocation policies

Host groups are optional. If you do not choose to set them up, you can mount file systems on clients and tune host settings and allocation policies on an individual level.

How host groups work

In the simplest case, the host groups functionality allows you to perform an allowed operation on all StoreAll clients by executing a command on the default `clients` host group with the CLI or the GUI. The `clients` host group includes all StoreAll clients configured in the cluster.

NOTE: The command intention is stored on the Fusion Manager until the next time the clients contact the Fusion Manager. (To force this contact, restart StoreAll software services on the clients, reboot the clients, or execute `ibrix_lwmount -a` or `ibrix_lwhost --a`.) When contacted, the Fusion Manager informs the clients about commands that were executed on host groups to which they belong. The clients then use this information to perform the operation.

You can also use host groups to perform different operations on different sets of clients. To do this, create a *host group tree* that includes the necessary host groups. You can then assign the clients manually, or the Fusion Manager can automatically perform the assignment when you register a StoreAll client, based on the client's cluster subnet. To use automatic assignment, create a domain rule that specifies the cluster subnet for the host group.

Creating a host group tree

The `clients` host group is the root element of the host group tree. Each host group in a tree can have only one parent, but a parent can have multiple children. In a host group tree, operations performed on lower-level nodes take precedence over operations performed on higher-level nodes. This means that you can effectively establish global client settings that you can override for specific clients.

For example, suppose that you want all clients to be able to mount file system `ifs1` and to implement a set of host tunings denoted as Tuning 1, but you want to override these global settings for certain host groups. To do this, mount `ifs1` on the `clients` host group, `ifs2` on host group A, `ifs3` on host group C, and `ifs4` on host group D, in any order. Then, set Tuning 1 on the `clients` host group and Tuning 2 on host group B. The end result is that all clients in host group B will mount `ifs1` and implement Tuning 2. The clients in host group A will mount `ifs2` and implement Tuning 1. The clients in host groups C and D respectively, will mount `ifs3` and `ifs4` and implement Tuning 1. The following diagram shows an example of these settings in a host group tree.

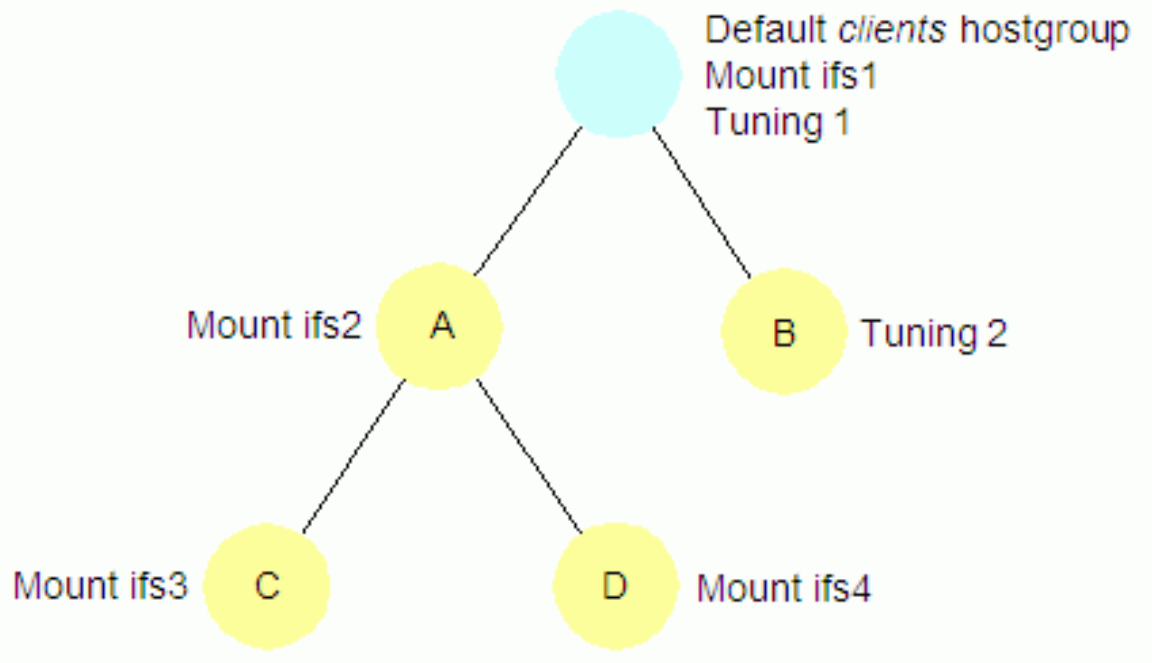

To create one level of host groups beneath the root, simply create the new host groups. You do not need to declare that the root node is the parent. To create lower levels of host groups, declare a parent element for host groups. Do not use a host name as a group name.

To create a host group tree using the CLI:

1. Create the first level of the tree:

   ```
   ibrix_hostgroup -c -g GROUPNAME
   ```

2. Create all other levels by specifying a parent for the group:

   ```
   ibrix_hostgroup -c -g GROUPNAME  [-p PARENT]
   ```

Adding a StoreAll client to a host group

You can add a StoreAll client to a host group or move a client to a different host group. All clients belong to the default `clients` host group.

To add or move a host to a host group, use the `ibrix_hostgroup` command as follows:

```
ibrix_hostgroup -m -g GROUP -h MEMBER
```

For example, to add the specified host to the `finance` group:

```
ibrix_hostgroup -m -g finance -h cl01.hp.com
```

Adding a domain rule to a host group

To configure automatic host group assignments, define a *domain rule* for host groups. A domain rule restricts host group membership to clients on a particular cluster subnet. The Fusion Manager uses the IP address that you specify for clients when you register them to perform a subnet match and sorts the clients into host groups based on the domain rules.

Setting domain rules on host groups provides a convenient way to centrally manage mounting, tuning, allocation policies, and preferred networks on different subnets of clients. A domain rule is a subnet IP address that corresponds to a client network. Adding a domain rule to a host group restricts its members to StoreAll clients that are on the specified subnet. You can add a domain rule at any time.

To add a domain rule to a host group, use the `ibrix_hostgroup` command as follows:

```
ibrix_hostgroup -a -g GROUPNAME -D DOMAIN
```

For example, to add the domain rule `192.168` to the `finance` group:

```
ibrix_hostgroup -a -g finance -D 192.168
```

Viewing host groups

To view all host groups or a specific host group, use the following command:

```
ibrix_hostgroup -l [-g GROUP]
```

Deleting host groups

When you delete a host group, its members are reassigned to the parent of the deleted group.

To force the reassigned StoreAll clients to implement the mounts, tunings, network interface preferences, and allocation policies that have been set on their new host group, either restart StoreAll software services on the clients or execute the following commands locally:

- `ibrix_lwmount -a` to force the client to pick up mounts or allocation policies
- `ibrix_lwhost --a` to force the client to pick up host tunings

To delete a host group using the CLI:

```
ibrix_hostgroup -d -g GROUPNAME
```

Other host group operations

Additional host group operations are described in the following locations:

- Creating or deleting a mountpoint, and mounting or unmounting a file system (see "Creating and mounting file systems" in the *HP StoreAll Storage File System User Guide*)
- Changing host tuning parameters (see "Tuning file serving nodes and StoreAll clients" (page 110))
- Preferring a network interface (see "Preferring network interfaces" (page 123))
- Setting allocation policy (see "Using file allocation" in the *HP StoreAll Storage File System User Guide*)

9 Monitoring cluster operations

This chapter describes how to monitor the operational state of the cluster and how to monitor cluster health.

Monitoring 9300/9320 hardware

The GUI displays status, firmware versions, and device information for the servers, virtual chassis, and system storage included in 9300 and 9320 systems.

Monitoring servers

To view information about the server and chassis included in your system.

1. Select **Servers** from the Navigator tree.

   The Servers panel lists the servers included in each chassis.

2. Select the server you want to obtain more information about.

   Information about the servers in the chassis is displayed in the right pane.

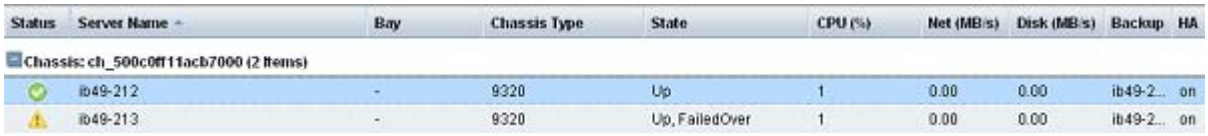

Status	Server Name	Bay	Chassis Type	State	CPU (%)	Net (MB/s)	Disk (MB/s)	Backup	HA
Chassis: ch_500c0ff11acb7000 (2 items)									
	ib49-212	-	9320	Up	1	0.00	0.00	ib49-2...	on
	ib49-213	-	9320	Up, FailedOver	1	0.00	0.00	ib49-2...	on

To view summary information for the selected server, select the **Summary** node in the lower Navigator tree.

Summary

Name	Value
State	Up
Group	servers
Standby	ib128-16s2
Auto Failover Enabled	Yes
Module	Loaded
ID	27f1e77f-b668-43fb-9c25-279cb3e783b8
Uptime	13 days, 19:23
Last Update	Thu Oct 11 17:24:26 UTC 2012
Admin IP	10.28.16.1
Filesystem Version	6.2.332(internal rev 130489 in SVN)
IAD Version	6.2.332
Protocol	TCP
OS	GNU/Linux
Kernel Version	2.6.18-194.el5
Architecture	x86_64
Processor	x86_64
CPU(System,User,Util,Nice)	0, 1, 2, 0
Load(1,5,15 min)	0.27, 0.32, 0.22
Memory Total (MB)	48,169.57
Memory Free (MB)	45,664.70
Buffers (KB)	402844
Cached (KB)	758500
Swap Total (MB)	16,383.74
Swap Free (MB)	16,383.74
Network (MB/s)	0.00
Disk (MB/s)	0.00
Number of Admin Threads	10
Number of Server Threads	128

Select the server component that you want to view from the lower Navigator panel, such as NICs.

The following are the top-level options provided for the server:

NOTE: Information about the **Hardware** node can be found in "Monitoring hardware components" (page 88).

- **HBAs**. The HBAs panel displays the following information:
  - Node WWN
  - Port WWN
  - Backup

- Monitoring
- State

- **NICs**. The NICs panel shows all NICs on the server, including offline NICs. The NICs panel displays the following information:
  - Name
  - IP
  - Type
  - State
  - Route
  - Standby Server
  - Standby Interface
- **Mountpoints**. The Mountpoints panel displays the following information:
  - Mountpoint
  - Filesystem
  - Access
- **NFS**. The NFS panel displays the following information:
  - Host
  - Path
  - Options
- **CIFS**. The CIFS panel displays the following information:

  NOTE: CIFS in the GUI has not been rebranded to SMB yet. CIFS is just a different name for SMB.

  - Name
  - Value
- **Power**. The Power panel displays the following information:
  - Host
  - Name
  - Type
  - IP Address
  - Slot ID

- **Events**. The Events panel displays the following information:
  - Level
  - Time
  - Event
- **Hardware**. The Hardware panel displays the following information:
  - The name of the hardware component.
  - The information gathered in regards to that hardware component.

  See "Monitoring hardware components" (page 88) for detailed information about the Hardware panel.

Monitoring hardware components

The Management Console provides information about the server hardware and its components. The 9300/9320 can be grouped virtually, which the Management Console interprets as a virtual chassis.

To monitor these components from the GUI:

1. Click **Servers** from the upper Navigator tree.
2. Click **Hardware** from the lower Navigator tree for information about the chassis that contains the server selected on the Servers panel, as shown in the following image.

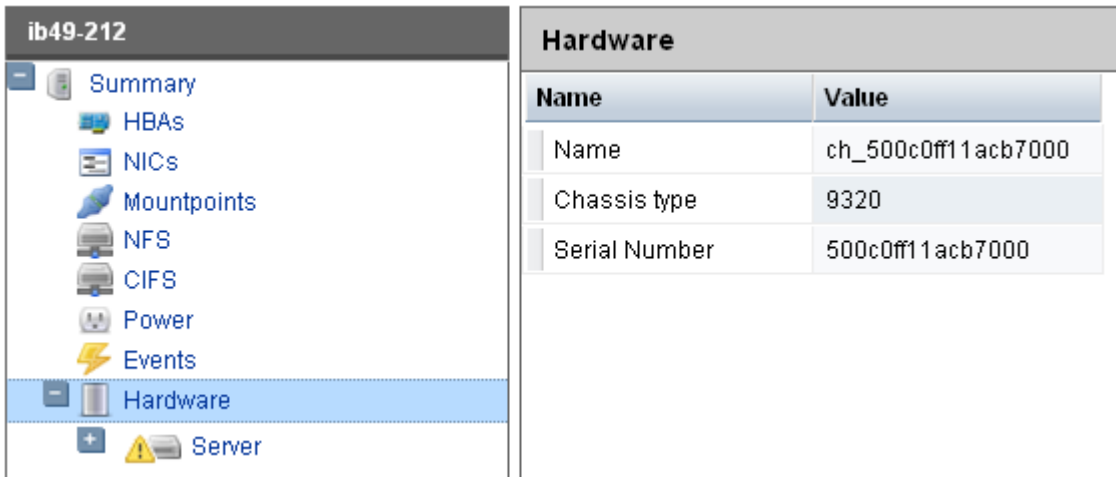

Obtaining server details

The Management Console provides detailed information for each server in the chassis. To obtain summary information for a server, select the **Server** node under the **Hardware** node.

The following overview information is provided for each server:

- Status
- Type
- Name
- UUID
- Serial number
- Model
- Firmware version

- Message[1]
- Diagnostic Message[1]

[1]Column dynamically appears depending on the situation.

Obtain detailed information for hardware components in the server by clicking the nodes under the Server node.

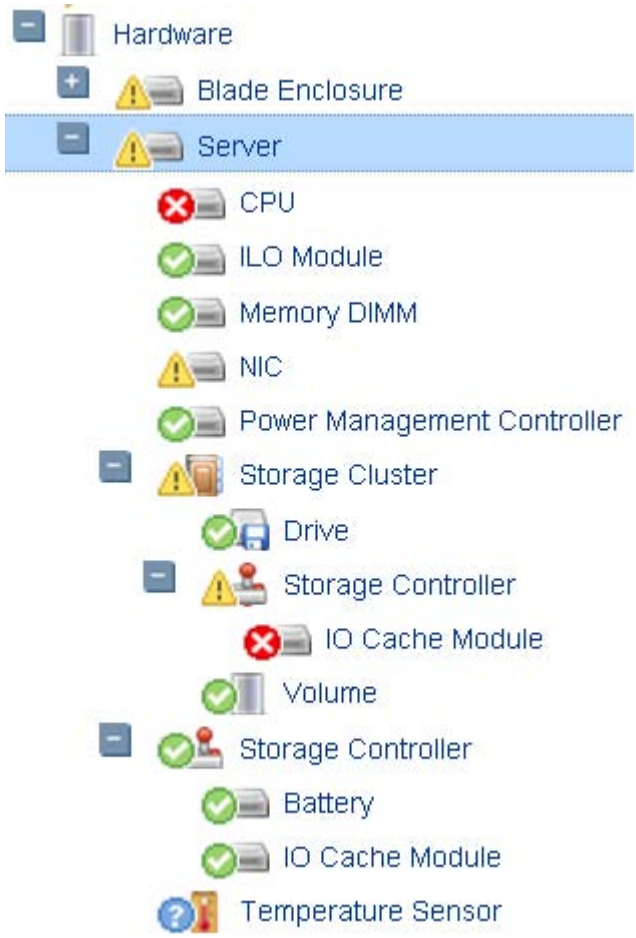

Table 2 Obtaining detailed information about a server

Panel name	Information provided
CPU	• Status • Type • Name • UUID • Model • Location
ILO Module	• Status • Type • Name • UUID • Serial Number • Model • Firmware Version • Properties
Memory DiMM	• Status • Type • Name • UUID • Location • Properties
NIC	• Status • Type • Name • UUID • Properties
Power Management Controller	• Status • Type • Name • UUID • Firmware Version
Storage Cluster	• Status • Type • Name • UUID

Table 2 Obtaining detailed information about a server *(continued)*

Panel name	Information provided
Drive: Displays information about each drive in a storage cluster.	• Status • Type • Name • UUID • Serial Number • Model • Firmware Version • Location • Properties
Storage Controller (Displayed for a server)	• Status • Type • Name • UUID • Serial Number • Model • Firmware Version • Location • Message • Diagnostic message
Volume: Displays volume information for each server.	• Status • Type • Name • UUID • Properties
Storage Controller (Displayed for a storage cluster)	• Status • Type • UUID • Serial Number • Model • Firmware Version • Message • Diagnostic Message
Battery (Displayed for each storage controller)	• Status • Type • UUID • Properties
IO Cache Module (Displayed for a storage controller)	• Status • Type • UUID • Properties

Table 2 Obtaining detailed information about a server *(continued)*

Panel name	Information provided
Temperature Sensor: Displays information for each temperature sensor.	• Status • Type • Name • UUID • Locations • Properties

Monitoring storage and storage components

Select **Vendor Storage** from the Navigator tree to display status and device information for storage and storage components.

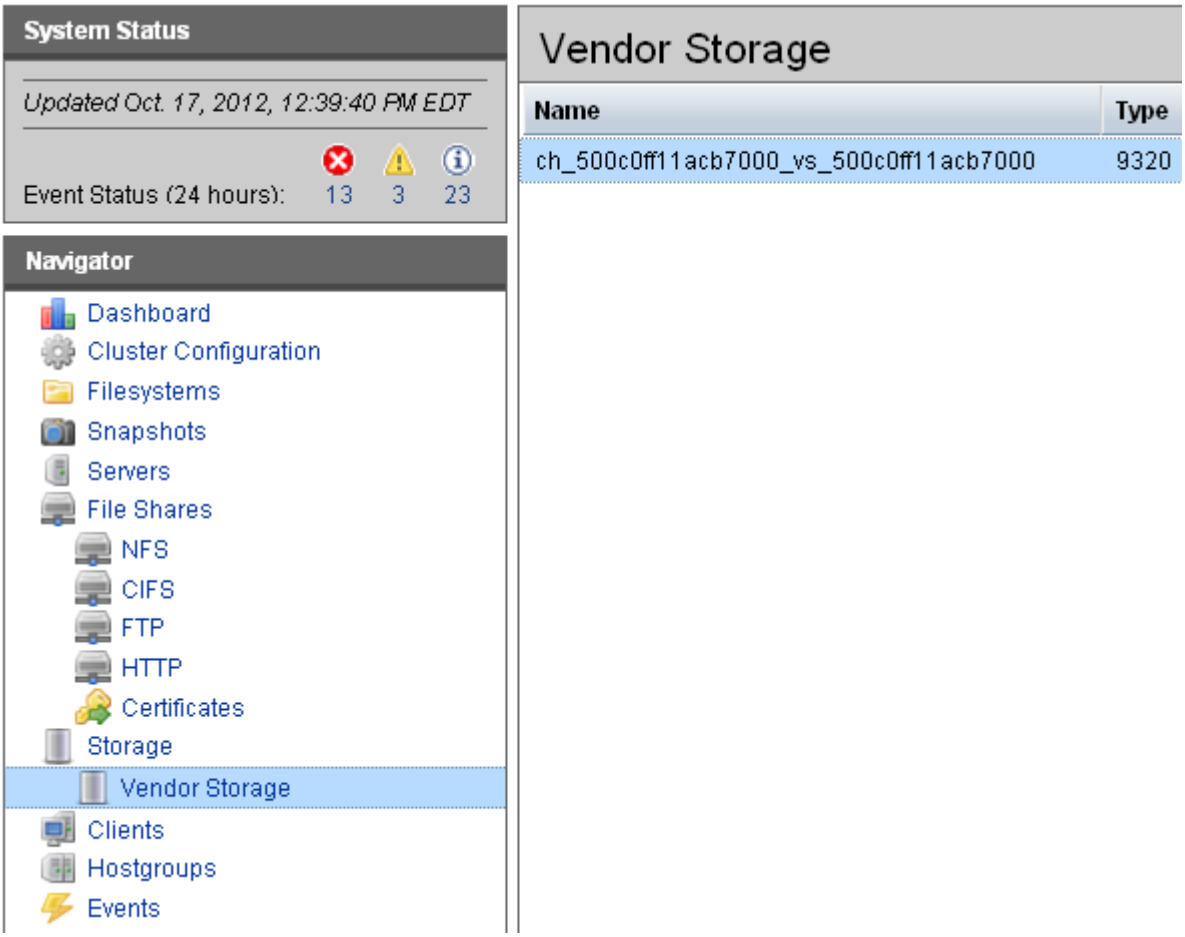

The Summary panel shows details for a selected vendor storage, as shown in the following image:

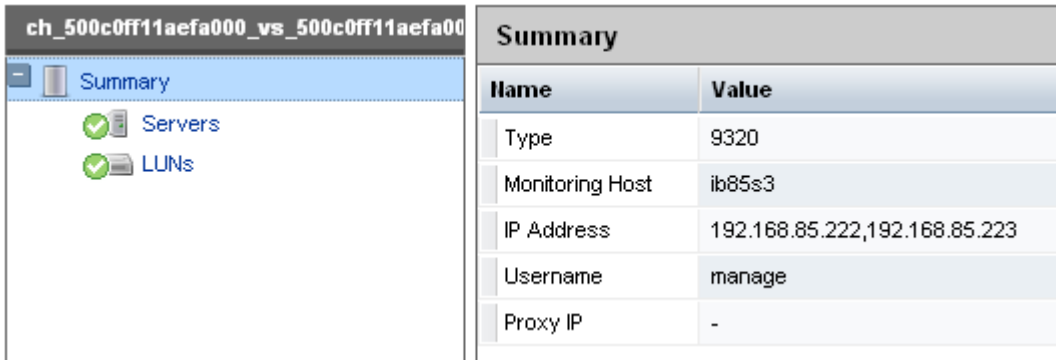

The Management Console provides a wide-range of information in regards to vendor storage.

Drill down into the following components in the lower Navigator tree to obtain additional details:

- **Servers**. The Servers panel lists the host names for the attached storage.
- **LUNs**. The LUNs panel provides information about the LUNs in a storage cluster. See "Managing LUNs in a storage cluster" (page 93) for more information.

Managing LUNs in a storage cluster

The LUNs panel provides information about the LUNs in a storage cluster. The following information is provided in the LUNs panel:

- LUN ID
- Physical Volume Name
- Physical Volume UUID

In the following image, the LUNs panel displays the LUNs for a storage cluster.

LUNs

Volume Name	LUN ID	Physical Volume Name	Physical Volume UUID
fs2_seg1	00c0ff11705200005fade...	d9	PdgjOy-ae60-ANEU-TtP...
vd01_v001	00c0ff1170520000afdba...	d1	ANMg1B-17Xf-1JcO-R0F...
vd02_v001	00c0ff1089810000dcdb...	d2	pG4i9x-O8Nx-Z4Zc-yD5x...
vd03_v001	00c0ff1170520000fcdba...	d3	FD0p8K-EVHt-53FM-Njj...
vd04_v001	00c0ff10898100001cdc...	d4	xxFJYP-LFZg-ShMP-Cqy...
vd05_v001	00c0ff117052000044dc...	d5	aHVEV5-3lp0-HyCc-GR...
vd06_v001	00c0ff108981000062dc...	d6	eU4h6c-rGwP-GHCU-zT...
vd07_v001	00c0ff117052000078dc...	d7	mUHuAF-1Uhc-vwrB-2s...
vd08_v001	00c0ff10898100008edc...	d8	y4hi3D-ri3L-6mrg-Fsd0-...

Monitoring the status of file serving nodes

The dashboard on the GUI displays information about the operational status of file serving nodes, including CPU, I/O, and network performance information.

To view this information from the CLI, use the `ibrix_server -l` command, as shown in the following sample output:

```
ibrix_server -l

SERVER_NAME   STATE          CPU(%)   NET_IO(MB/s)   DISK_IO(MB/s)   BACKUP   HA
```

```
-----------  ------------  ------  ------------  -------------  ------  --
node1        Up, HBAsDown       0          0.00           0.00          off
node2        Up, HBAsDown       0          0.00           0.00          off
```

File serving nodes can be in one of three operational states: Normal, Alert, or Error. These states are further broken down into categories describing the failover status of the node and the status of monitored NICs and HBAs.

State	Description
Normal	**Up:** Operational.
Alert	**Up-Alert:** Server has encountered a condition that has been logged. An event will appear in the Status tab of the GUI, and an email notification may be sent. **Up-InFailover:** Server is powered on and visible to the Fusion Manager, and the Fusion Manager is failing over the server's segments to a standby server. **Up-FailedOver:** Server is powered on and visible to the Fusion Manager, and failover is complete.
Error	**Down-InFailover:** Server is powered down or inaccessible to the Fusion Manager, and the Fusion Manager is failing over the server's segments to a standby server. **Down-FailedOver:** Server is powered down or inaccessible to the Fusion Manager, and failover is complete. **Down:** Server is powered down or inaccessible to the Fusion Manager, and no standby server is providing access to the server's segments.

The STATE field also reports the status of monitored NICs and HBAs. If you have multiple HBAs and NICs and some of them are down, the state is reported as HBAsDown or NicsDown.

Monitoring cluster events

StoreAll software events are assigned to one of the following categories, based on the level of severity:

- **Alerts.** A disruptive event that can result in loss of access to file system data. For example, a segment is unavailable or a server is unreachable.
- **Warnings.** A potentially disruptive condition where file system access is not lost, but if the situation is not addressed, it can escalate to an alert condition. Some examples are reaching a very high server CPU utilization or nearing a quota limit.
- **Information.** An event that changes the cluster (such as creating a segment or mounting a file system) but occurs under normal or nonthreatening conditions.

Events are written to an events table in the configuration database as they are generated. To maintain the size of the file, HP recommends that you periodically remove the oldest events. See "Removing events from the events database table" (page 95).

You can set up event notifications through email (see "Setting up email notification of cluster events" (page 70)) or SNMP traps (see "Setting up SNMP notifications" (page 72)).

Viewing events

The GUI dashboard specifies the number of events that have occurred in the last 24 hours. Click **Events** in the GUI Navigator to view a report of the events. You can also view events that have been reported for specific file systems or servers.

On the CLI, use the `ibrix_event` command to view information about cluster events.

To view events by alert type, use the following command:

```
ibrix_event -q [-e ALERT|WARN|INFO]
```

The `ibrix_event -l` command displays events in a short format; event descriptions are truncated to fit on one line. The `-n` option specifies the number of events to display. The default is 100.

```
$ ibrix_event -l -n 3
EVENT ID  TIMESTAMP         LEVEL  TEXT
--------  ---------------   -----  ----
    1983  Feb 14 15:08:15   INFO   File system ifs1 created
    1982  Feb 14 15:08:15   INFO   Nic eth0[99.224.24.03] on host ix24-03.ad.hp.com up
    1981  Feb 14 15:08:15   INFO   Ibrix kernel file system is up on ix24-03.ad.hp.com
```

The `ibrix_event -i` command displays events in long format, including the complete event description.

```
$ ibrix_event -i -n 2
Event:
=======
EVENT ID        :  1981
TIMESTAMP       :  Feb 14 15:08:15
LEVEL           :  INFO
TEXT            :  Ibrix kernel file system is up on ix24-03.ad.hp.com
FILESYSTEM      :
HOST            :  ix24-03.ad.hp.com
USER NAME       :
OPERATION       :
SEGMENT NUMBER  :
PV NUMBER       :
NIC             :
HBA             :
RELATED EVENT   :  0

Event:
=======
EVENT ID        :  1980
TIMESTAMP       :  Feb 14 15:08:14
LEVEL           :  ALERT
TEXT            :  category:CHASSIS, name: 9730_ch1, overallStatus:DEGRADED,
component:OAmodule, uuid:09USE038187WOAModule2, status:MISSING, Message: The Onboard
Administrator module is missing or has failed., Diagnostic message: Reseat the Onboard
Administrator module. If reseating the module does not resolve the issue, replace the Onboard
Administrator module., eventId:000D0004, location:OAmodule in chassis S/N:USE123456W,
level:ALERT
FILESYSTEM      :
HOST            :  ix24-03.ad.hp.com
USER NAME       :
OPERATION       :
SEGMENT NUMBER  :
PV NUMBER       :
NIC             :
HBA             :
RELATED EVENT   :  0
```

The `ibrix_event -l` and `-i` commands can include options that act as filters to return records associated with a specific file system, server, alert level, and start or end time. See the *HP StoreAll Network Storage System CLI Reference Guide* for more information.

Removing events from the events database table

Use the `ibrix_event -p` command to removes event from the events table, starting with the oldest events. The default is to remove the oldest seven days of events. To change the number of days, include the `-o` *`DAYS_COUNT`* option.

```
ibrix_event -p [-o DAYS_COUNT]
```

Monitoring cluster health

To monitor the functional health of file serving nodes and StoreAll clients, execute the `ibrix_health` command. This command checks host performance in several functional areas and provides either a summary or a detailed report of the results.

Health checks

The `ibrix_health` command runs these health checks on file serving nodes:

- Pings remote file serving nodes that share a network with the test hosts. Remote servers that are pingable might not be connected to a test host because of a Linux or StoreAll software issue. Remote servers that are not pingable might be down or have a network problem.
- If test hosts are assigned to be network interface monitors, pings their monitored interfaces to assess the health of the connection. (For information on network interface monitoring, see "Setting network interface options in the configuration database" (page 123)
- Determines whether specified hosts can read their physical volumes.

The `ibrix_health` command runs this health check on both file serving nodes and StoreAll clients:

- Determines whether information maps on the tested hosts are consistent with the configuration database.

If you include the `-b` option, the command also checks the health of standby servers (if configured).

Health check reports

The summary report provides an overall health check result for all tested file serving nodes and StoreAll clients, followed by individual results. If you include the `-b` option, the standby servers for all tested file serving nodes are included when the overall result is determined. The results will be one of the following:

- **Passed.** All tested hosts and standby servers passed every health check.
- **Failed.** One or more tested hosts failed a health check. The health status of standby servers is not included when this result is calculated.
- **Warning.** A suboptimal condition that might require your attention was found on one or more tested hosts or standby servers.

The detailed report consists of the summary report and the following additional data:

- Summary of the test results
- Host information such as operational state, performance data, and version data
- Nondefault host tunings
- Results of the health checks

By default, the Result Information field in a detailed report provides data only for health checks that received a Failed or a Warning result. Optionally, you can expand a detailed report to provide data about checks that received a Passed result, as well as details about the file system and segments.

Viewing a summary health report

To view a summary health report, use the `ibrix_health -l` command:

```
ibrix_health -l [-h HOSTLIST] [-f] [-b]
```

By default, the command reports on all hosts. To view specific hosts, include the `-h HOSTLIST` argument. To view results only for hosts that failed the check, include the `-f` argument. To include standby servers in the health check, include the `-b` argument.

The following is an example of the output from the `ibrix_health -l` command:

```
[root@bv18-03 ~]# ibrix_health -l
 Overall Health Checker Results - PASSED
 =====================================
 Host Summary Results
 ====================
```

```
  Host     Result  Type    State  Network     Last Update
  -------  ------  ------  -----  ----------  -----------
  bv18-03  PASSED  Server  Up     10.10.18.3  Thu Oct 25 14:23:12 MDT 2012
  bv18-04  PASSED  Server  Up     10.10.18.4  Thu Oct 25 14:23:22 MDT 2012
```

Viewing a detailed health report

To view a detailed health report, use the `ibrix_health -i` command:

`ibrix_health -i -h` *`HOSTLIST`* `[-f] [-s] [-v]`

The `-f` option displays results only for hosts that failed the check. The `-s` option includes information about the file system and its segments. The `-v` option includes details about checks that received a Passed or Warning result.

The following example shows a detailed health report for file serving node `bv18-04`:

```
[root@bv18-04 ~]# ibrix_health -i -h bv18-04
  Overall Health Checker Results - PASSED
  =======================================
  Host Summary Results
  ====================
  Host     Result  Type    State  Network     Last Update
  -------  ------  ------  -----  ----------  -----------
  bv18-04  PASSED  Server  Up     10.10.18.4  Thu Oct 25 13:59:40 MDT 2012
    Report
    ======
    Overall Result
    ==============
    Result  Type    State  Module  Up time    Last Update                  Network     Thread  Protocol
    ------  ------  -----  ------  ---------  ---------------------------  ----------  ------  --------
    PASSED  Server  Up     Loaded  1699630.0  Thu Oct 25 13:59:40 MDT 2012  10.10.18.4  64      true
    CPU Information
    ===============
    Cpu(System,User,Util,Nice)  Load(1,3,15 min)  Network(Bps)  Disk(Bps)
    --------------------------  ----------------  ------------  ---------
    0, 0, 0, 0                  0.09, 0.05, 0.01          1295       1024
    Memory Information
    ==================
    Mem Total  Mem Free  Buffers(KB)  Cached(KB)  Swap Total(KB)  Swap Free(KB)
    ---------  --------  -----------  ----------  --------------  -------------
      8045992   4190584       243312     2858364        14352376       14352376
    Version/OS Information
    ======================
    Fs Version   IAD Version  OS   OS Version                    Kernel Version  Architecture  Processor
    ----------  -----------  -----------  -------------------------------------------------------------
---  ------------  ---------
    6.3.72   6.3.72      GNU/Linux  Red Hat Enterprise Linux Server release 5.5 (Tikanga)   2.6.18-194.
el5  x86_64        x86_64
    Remote Hosts
    ============
    Host     Type    Network     Protocol  Connection State
    -------  ------  ----------  --------  ----------------
    bv18-03  Server  10.10.18.3  true       S_SET S_READY S_SENDHB
    bv18-04  Server  10.10.18.4  true       S_NEW
      Check Results
      =============
      Check : bv18-04 can ping remote segment server hosts
      ====================================================
      Check Description               Result  Result Information
      ------------------------------  ------  ------------------
      Remote server bv18-03 pingable  PASSED
      Check : Iad's monitored nics are pingable
      =========================================
      Check Description                                  Result  Result Information
      -------------------------------------------------  ------  ------------------
      User nic bv18-04/bond1:2 pingable from host bv18-03  PASSED
      Check : Physical volumes are readable
      =====================================
      Check Description                                                  Result  Result Information
      -----------------------------------------------------------------  ------  ------------------
      Physical volume 0wndzX-STuL-wSIi-wc7w-12hv-JZ2g-Lj2JTf readable  PASSED  /dev/mpath/mpath2
      Physical volume aoA402-I1ek-G9B2-HHyR-H5Y8-eexU-P6knhd readable  PASSED  /dev/mpath/mpath1
      Physical volume h7krR6-2pxA-M8bD-dkdf-3PK7-iwFE-L17jcD readable  PASSED  /dev/mpath/mpath0
      Physical volume voXTso-a2KQ-MWCN-tGcu-1OBs-ejWG-YrKLEe readable  PASSED  /dev/mpath/mpath3
      Check : Iad and Fusion Manager consistent
      =========================================
      Check Description
Result  Result Information
   ------------  ------  ------------------
      bv18-03 engine uuid matches on Iad and Fusion Manager                    PASSED
      bv18-03 IP address matches on Iad and Fusion Manager                     PASSED
      bv18-03 network protocol matches on Iad and Fusion Manager               PASSED
      bv18-03 engine connection state on Iad is up                             PASSED
      bv18-04 engine uuid matches on Iad and Fusion Manager                    PASSED
      bv18-04 IP address matches on Iad and Fusion Manager                     PASSED
      bv18-04 network protocol matches on Iad and Fusion Manager               PASSED
      bv18-04 engine connection state on Iad is up                             PASSED
```

```
ibrixFS file system uuid matches on Iad and Fusion Manager                PASSED
ibrixFS file system generation matches on Iad and Fusion Manager          PASSED
ibrixFS file system number segments matches on Iad and Fusion Manager     PASSED
ibrixFS file system mounted state matches on Iad and Fusion Manager       PASSED
Superblock owner for segment 4 of filesystem ibrixFS on bv18-04 matches
on Iad and Fusion Manager                                                 PASSED
Superblock owner for segment 3 of filesystem ibrixFS on bv18-04 matches
on Iad and Fusion Manager                                                 PASSED
Superblock owner for segment 2 of filesystem ibrixFS on bv18-04 matches
on Iad and Fusion Manager                                                 PASSED
Superblock owner for segment 1 of filesystem ibrixFS on bv18-04 matches
on Iad and Fusion Manager                                                 PASSED
```

Viewing logs

Logs are provided for the Fusion Manager, file serving nodes, and StoreAll clients. Contact HP Support for assistance in interpreting log files. You might be asked to tar the logs and email them to HP.

Viewing operating statistics for file serving nodes

Periodically, the file serving nodes report the following statistics to the Fusion Manager:

- **Summary.** General operational statistics including CPU usage, disk throughput, network throughput, and operational state. For information about the operational states, see "Monitoring the status of file serving nodes" (page 93).
- **IO.** Aggregate statistics about reads and writes.
- **Network.** Aggregate statistics about network inputs and outputs.
- **Memory.** Statistics about available total, free, and swap memory.
- **CPU.** Statistics about processor and CPU activity.
- **NFS.** Statistics about NFS client and server activity.

The GUI displays most of these statistics on the dashboard. See "Using the StoreAll Management Console" (page 29) for more information.

To view the statistics from the CLI, use the following command:

```
ibrix_stats -l [-s] [-c] [-m] [-i] [-n] [-f] [-h HOSTLIST]
```

Use the options to view only certain statistics or to view statistics for specific file serving nodes:

`-s` Summary statistics
`-c` CPU statistics
`-m` Memory statistics
`-i` I/O statistics
`-n` Network statistics
`-f` NFS statistics
`-h` The file serving nodes to be included in the report

Sample output follows:

```
---------Summary------------
HOST                Status  CPU Disk(MB/s) Net(MB/s)
lab12-10.hp.com     Up      0   22528      616
---------IO-------------
HOST Read(MB/s)     Read(IO/s) Read(ms/op) Write(MB/s) Write(IO/s) Write(ms/op)
lab12-10.hp.com    22528      2           5           0           0.00
---------Net------------
HOST                In(MB/s) In(IO/s) Out(MB/s) Out(IO/s)
lab12-10.hp.com     261      3        355       2
---------Mem------------
HOST                MemTotal(MB) MemFree(MB) SwapTotal(MB) SwapFree(MB)
lab12-10.hp.com     1034616      703672      2031608       2031360
---------CPU------------
HOST                User System Nice Idle IoWait Irq SoftIrq
lab12-10.hp.com     0    0      0    0    97     1   0
```

```
---------NFS v3--------
HOST                 Null Getattr Setattr Lookup Access Readlink Read Write
lab12-10.hp.com       0    0       0       0      0      0        0    0

HOST                Create Mkdir Symlink Mknod Remove Rmdir Rename
lab12-10.hp.com     0      0     0       0     0      0     0

HOST                Link Readdir Readdirplus Fsstat Fsinfo Pathconf Commit
lab12-10.hp.com     0    0       0           0       0     0        0
```

10 Using the Statistics tool

The Statistics tool reports historical performance data for the cluster or for an individual file serving node. You can view data for the network, the operating system, and the file systems, including the data for NFS, memory, and block devices. Statistical data is transmitted from each file serving node to the Fusion Manager, which controls processing and report generation.

Installing and configuring the Statistics tool

The Statistics tool has two main processes:

- **Manager process.** This process runs on the active Fusion Manager. It collects and aggregates cluster-wide statistics from file serving nodes running the Agent process, and also collects local statistics. The Manager generates reports based on the aggregated statistics and collects reports from all file serving nodes. The Manager also controls starting and stopping the Agent process.
- **Agent process.** This process runs on the file serving nodes. It collects and aggregates statistics on the local system and generates reports from those statistics.

IMPORTANT: The Statistics tool uses remote file copy (`rsync`) to move statistics data from the file serving nodes to the Fusion Manager for processing, report generation, and display. SSH keys are configured automatically across all the file serving nodes to the active Fusion Manager.

Installing the Statistics tool

The Statistics tool is installed automatically when the StoreAll software is installed on the file serving nodes. To install or reinstall the Statistics tool manually, use the following command:

```
ibrixinit -tt
```

Note the following:

- Installation logs are located at `/tmp/stats-install.log`.
- By default, installing the Statistics tool does not start the Statistics tool processes. See "Controlling Statistics tool processes" (page 105) for information about starting and stopping the processes.
- If the Fusion Manager deamon is not running during the installation, Statstool is installed as passive. When Fusion Manager acquires an active/passive state, the Statstool management console automatically changes according to the state of Fusion Manager.

Enabling collection and synchronization

To enable collection and synchronization, configure synchronization between nodes. Run the following command on the active Fusion Manager node, specifying the node names of **all** file serving nodes:

```
/usr/local/ibrix/stats/bin/stmanage setrsync <node1_name> ...
<nodeN_name>
```

For example:

```
# stmanage setrsync ibr-3-31-1 ibr-3-31-2 ibr-3-31-3
```

NOTE: Do not run the command on individual nodes. All nodes must be specified in the same command and can be specified in any order. Be sure to use node names, not IP addresses.

To test the `rsync` mechanism, see "Testing access" (page 105).

Upgrading the Statistics tool from StoreAll software 6.0

The statistics history is retained when you upgrade to version 6.1 or later.

The Statstool software is upgraded when the StoreAll software is upgraded using the `ibrix_upgrade` and `auto_ibrixupgrade` scripts.

Note the following:

- If statistics processes were running before the upgrade started, those processes will automatically restart after the upgrade completes successfully. If processes were not running before the upgrade started, you must start them manually after the upgrade completes.
- If the Statistics tool was not previously installed, the StoreAll software upgrade installs the tool but the Statistic processes are not started. For information about starting the processes, see"Controlling Statistics tool processes" (page 105).
- Configurable parameters (such as `age.retain.files=24h`) set in the `/etc/ibrix/stats.conf` file before the upgrade are not retained after the upgrade.
- After the upgrade, historical data and reports are moved from the `/var/lib/ibrix/histstats` folder to the `/local/statstool/histstats` folder.
- The upgrade retains the Statistics tool database but not the reports. You can regenerate reports for the data stored before the upgrade by specifying the date range. See "Generating reports" (page 102).

Using the Historical Reports GUI

You can use the GUI to view or generate reports for the entire cluster or for a specific file serving node. To open the GUI, select **Historical Reports** on the GUI dashboard.

NOTE: By default, installing the Statistics tool does not start the Statistics tool processes. The GUI displays a message if the processes are not running on the active Fusion Manager. (No message appears if the processes are already running on the active Fusion Manager, or if the processes are not running on any of the passive management consoles.) See "Controlling Statistics tool processes" (page 105) for information about starting the processes.

The statistics home page provides three views, or formats, for listing the reports. Following is the Simple View, which sorts the reports according to type (hourly, daily, weekly, detail).

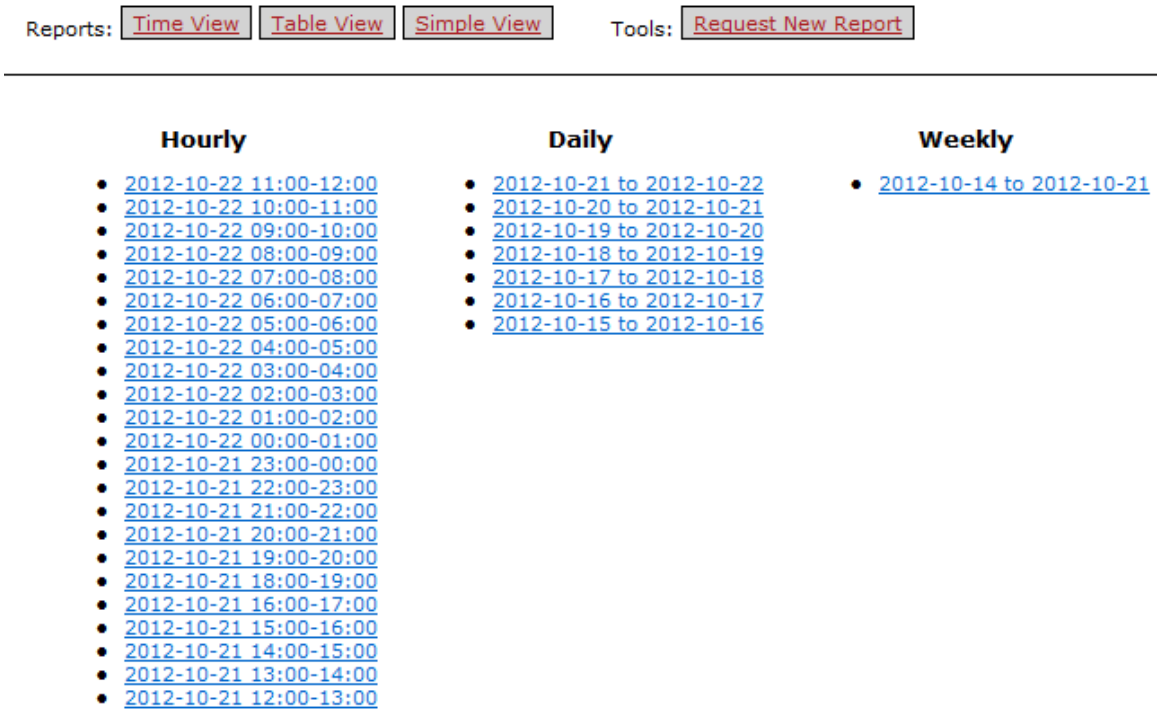

The Time View lists the reports in chronological order, and the Table View lists the reports by cluster or server. Click a report to view it.

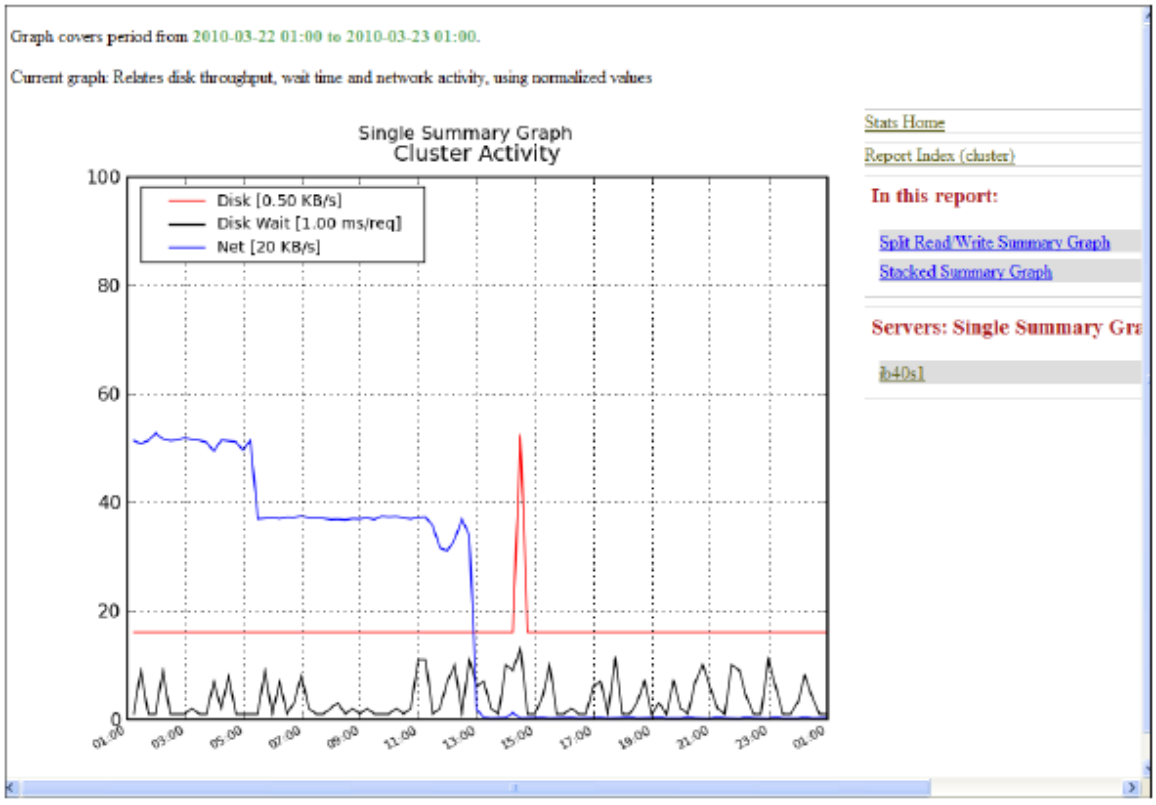

Generating reports

To generate a new report, click **Request New Report** on the StoreAll Management Console Historical Reports GUI.

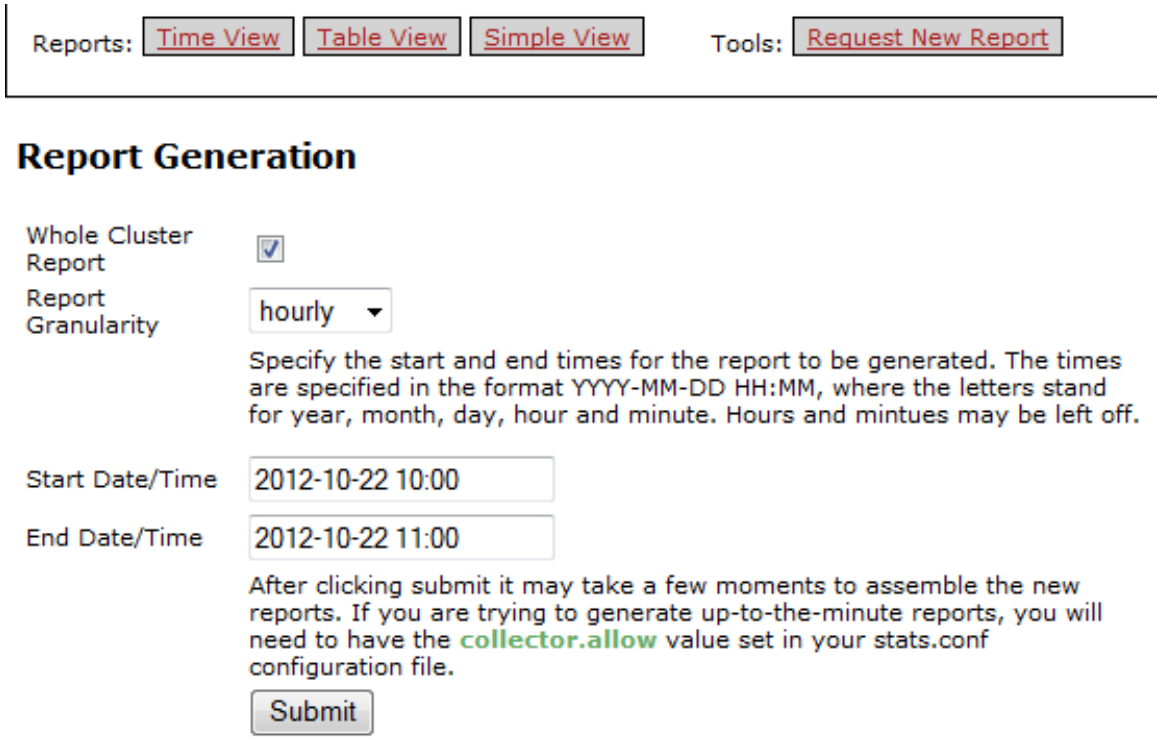
Reports:
Time View
Table View
Simple View
Tools:
Request New Report
Report Generation
Whole Cluster Report
Report Granularity
hourly
Specify the start and end times for the report to be generated. The times are specified in the format YYYY-MM-DD HH:MM, where the letters stand for year, month, day, hour and minute. Hours and mintues may be left off.
Start Date/Time
2012-10-22 10:00
End Date/Time
2012-10-22 11:00
After clicking submit it may take a few moments to assemble the new reports. If you are trying to generate up-to-the-minute reports, you will need to have the collector.allow value set in your stats.conf configuration file.
Submit

To generate a report, enter the necessary specifications and click **Submit**. The completed report appears in the list of reports on the statistics home page.

When generating reports, be aware of the following:

- A report can be generated only from statistics that have been gathered. For example, if you start the tool at 9:40 a.m. and ask for a report from 9:00 a.m. to 9:30 a.m., the report cannot be generated because data was not gathered for that period.
- Reports are generated on an hourly basis. It may take up to an hour before a report is generated and made available for viewing.

NOTE: If the system is currently generating reports and you request a new report at the same time, the GUI issues an error. Wait a few moments and then request the report again.

Deleting reports

To delete a report, log into each node and remove the report from the `/local/statstool/histstats/reports/` directory.

Maintaining the Statistics tool

Space requirements

The Statistics tool requires about 4 MB per hour for a two-node cluster. To manage space, take the following steps:

- Maintain sufficient space (4 GB to 8 GB) for data collection in the `/usr/local/statstool/histstats` directory.
- Monitor the space in the `/local/statstool/histstats/reports/` directory. For the default values, see "Changing the Statistics tool configuration" (page 104).

Updating the Statistics tool configuration

When you first configure the Statistics tool, the configuration includes information for all file systems configured on the cluster. If you add a new node or a new file system, or make other additions to the cluster, you must update the Statistics tool configuration. Complete the following steps:

1. If you are adding a new file serving node to the cluster, enable synchronization for the node. See "Enabling collection and synchronization" (page 100) for more information.
2. Add the file system to the Statistics tool. Run the following command on the node hosting the active Fusion Manager:

   ```
   /usr/local/ibrix/stats/bin/stmanage loadfm
   ```

   The new configuration is updated automatically on the other nodes in the cluster. You do not need to restart the collection process; collection continues automatically.

Changing the Statistics tool configuration

You can change the configuration only on the management node. To change the configuration, add a configuration parameter and its value to the `/etc/ibrix/stats.conf` file on the currently active node. Do not modify the `/etc/ibrix/statstool.conf` and `/etc/ibrix/statstool.local.conf` files directly.

You can set the following parameters to specify the number of reports that are retained.

Parameter	Report Type to retain	Default Retention Period
`age.report.hourly`	Hourly report	1 day
`age.report.daily`	Daily report	7 days
`age.report.weekly`	Weekly report	14 days
`age.report.other`	User-generated report	7 days

For example, for daily reports, the default of 7 days saves seven reports. To save only three daily reports, set the `age.report.daily` parameter to 3 days:

```
age.report.daily=3d
```

NOTE: You do not need to restart processes after changing the configuration. The updated configuration is collected automatically.

Fusion Manager failover and the Statistics tool configuration

In a High Availability environment, the Statistics tool fails over automatically when the Fusion Manager fails over. You do not need to take any steps to perform the failover. The statistics configuration changes automatically as the Fusion Manager configuration changes.

The following actions occur after a successful failover:

- If Statstool processes were running before the failover, they are restarted. If the processes were not running, they are not restarted.
- The Statstool passive management console is installed on the StoreAll Fusion Manager in maintenance mode.
- `Setrsync` is run automatically on all cluster nodes from the current active Fusion Manager.
- `Loadfm` is run automatically to present all file system data in the cluster to the active Fusion Manager.
- The stored cluster-level database generated before the Fusion Manager failover is moved to the current active Fusion Manager, allowing you to request reports for the specified range if pre-generated reports are not available under the Hourly, Daily and Weekly categories. See "Generating reports" (page 102).

NOTE: If the old active Fusion Manager is not available (pingable) for more than two days, the historical statistics database is not transferred to the current active Fusion Manager.

- If configurable parameters were set before the failover, the parameters are retained after the failover.

Check the `/usr/local/ibrix/log/statstool/stats.log` for any errors.

NOTE: The reports generated before failover will not be available on the current active Fusion Manager.

Checking the status of Statistics tool processes

To determine the status of Statistics tool processes, run the following command:

```
#/etc/init.d/ibrix_statsmanager status
 ibrix_statsmanager (pid  25322) is running...
```

In the output, the `pid` is the process id of the "master" process.

Controlling Statistics tool processes

Statistics tool processes on all file serving nodes connected to the active Fusion Manager can be controlled remotely from the active Fusion Manager. Use the `ibrix_statscontrol` tool to start or stop the processes on all connected file serving nodes or on specified hostnames only.

- Stop processes on all file serving nodes, including the Fusion Manager:

  ```
  # /usr/local/ibrix/stats/bin/ibrix_statscontrol stopall
  ```

- Start processes on all file serving nodes, including the Fusion Manager:

  ```
  # /usr/local/ibrix/stats/bin/ibrix_statscontrol startall
  ```

- Stop processes on specific file serving nodes:

  ```
  # /usr/local/ibrix/stats/bin/ibrix_statscontrol stop <hostname1>
  <hostname2> ..
  ```

- Start processes on specific file serving nodes:

  ```
  # /usr/local/ibrix/stats/bin/ibrix_statscontrol start <hostname1>
  <hostname2> ..
  ```

Troubleshooting the Statistics tool

Testing access

To verify that `ssh` authentication is enabled and data can be obtained from the nodes without prompting for a password, run the following command:

```
# /usr/local/ibrix/stats/bin/stmanage testpull
```

Other conditions

- **Data is not collected.** If data is not being gathered in the common directory for the Statistics Manager (`/usr/local/statstool/histstats/` by default), restart the Statistics tool processes on all nodes. See "Controlling Statistics tool processes" (page 105).
- **Installation issues.** Check the `/tmp/stats-install.log` and try to fix the condition, or send the `/tmp/stats-install.log` to HP Support.
- **Missing reports for file serving nodes.** If reports are missing on the Stats tool web page, check the following:
  - Determine whether collection is enabled for the particular file serving node. If not, see "Enabling collection and synchronization" (page 100).
  - Check for time synchronization. All servers in the cluster should have the same date time and time zone to allow proper collection and viewing of reports.

Log files

See `/usr/local/ibrix/log/statstool/stats.log` for detailed logging for the Statistics tool. (The information includes detailed exceptions and traceback messages.) The logs are rolled over at midnight every day and only seven days of statistics logs are retained.

The default `/var/log/messages` log file also includes logging for the Statistics tool, but the messages are short.

Uninstalling the Statistics tool

The Statistics tool is uninstalled when the StoreAll software is uninstalled.

To uninstall the Statistics tool manually, use one of the following commands:

- Uninstall the Statistics tool, including the Statstics tool and dependency rpms:

  ```
  # ibrixinit -tt -u
  ```

- Uninstall the Statistics tool, retaining the Statstics tool and dependency rpms:

  ```
  # ibrixinit -tt -U
  ```

11 Maintaining the system

Shutting down the system

To shut down the system completely, first shut down the StoreAll software, and then power off the system hardware.

Shutting down the StoreAll software

Use the following procedure to shut down the StoreAll software. Unless noted otherwise, run the commands from the node hosting the active Fusion Manager.

1. Stop any active remote replication, data tiering, or rebalancer tasks. Run the following command to list active tasks and note their task IDs:

   ```
   # ibrix_task -l
   ```

   Run the following command to stop each active task, specifying its task ID:

   ```
   # ibrix_task -k -n TASKID
   ```

2. Disable High Availability on all cluster nodes:

   ```
   ibrix_server -m -U
   ```

3. Move all passive Fusion Manager instances into nofmfailover mode:

   ```
   ibrix_fm -m nofmfailover -A
   ```

4. Stop the SMB, NFS and NDMP services on all nodes. Run the following commands:

   ```
   ibrix_server -s -t cifs -c stop
   ibrix_server -s -t nfs -c stop
   ibrix_server -s -t ndmp -c stop
   ```

   If you are using SMB, verify that all likewise services are down on all file serving nodes:

   ```
   ps -ef | grep likewise
   ```

   Use `kill -9` to stop any likewise services that are still running.

   If you are using NFS, verify that all NFS processes are stopped:

   ```
   ps -ef | grep nfs
   ```

   If processes are running, use the following commands on the affected node:

   ```
   # pdsh -a service nfslock stop | dshbak
   # pdsh -a service nfs stop | dshbak
   ```

   If necessary, run the following command on all nodes to find any open file handles for the mounted file systems:

   ```
   lsof </mountpoint>
   ```

   Use `kill -9` to stop any processes that still have open file handles on the file systems.

5. List file systems mounted on the cluster:

   ```
   # ibrix_fs -l
   ```

6. Unmount all file systems from StoreAll clients:

   - On Linux StoreAll clients, run the following command to unmount each file system:

     ```
     ibrix_lwumount -f <fs_name>
     ```

   - On Windows StoreAll clients, stop all applications accessing the file systems, and then use the client GUI to unmount the file systems (for example, I: DRIVE). Next, go to Services and stop the fusion service.

7. Unmount all file systems on the cluster nodes:

   ```
   ibrix_umount -f <fs_name>
   ```

   To unmount file systems from the GUI, select **Filesystems** > **unmount**.

8. Verify that all file systems are unmounted:

   ```
   ibrix_fs -l
   ```

   If a file system fails to unmount on a particular node, continue with this procedure. The file system will be forcibly unmounted during the node shutdown.

9. Shut down all StoreAll Server services and verify the operation:

   ```
   # pdsh -a /etc/init.d/ibrix_server stop | dshbak
   # pdsh -a /etc/init.d/ibrix_server status | dshbak
   ```

10. Wait for the Fusion Manager to report that all file serving nodes are down:

    ```
    # ibrix_server -l
    ```

11. Shut down all nodes other than the node hosting the active Fusion Manager:

    ```
    # pdsh -w HOSTNAME shutdown -t now "now"
    ```

    For example:

    ```
    # pdsh -w x850s3 shutdown -t now "now"
    # pdsh -w x850s2 shutdown -t now "now"
    ```

12. Shut down the node hosting the active agile Fusion Manager:

    ```
    shutdown -t now “now”
    ```

13. Use `ping` to verify that the nodes are down. For example:

    ```
    # ping x850s2
    PING x850s2.l3domain.l3lab.com (12.12.80.102) 56(84) bytes of data.
    x850s1.l3domain.l3lab.com (12.12.82.101) icmp_seq=2 Destination Host
    Unreachable
    ```

    If you are unable to shut down a node cleanly, use the following command to power the node off using the iLO interface:

    ```
    # ibrix_server -P off -h HOSTNAME
    ```

14. Shut down the Fusion Manager services and verify:

    ```
    # /etc/init.d/ibrix_fusionmanager stop
    # /etc/init.d/ibrix_fusionmanager status
    ```

15. Shut down the node hosting the active Fusion Manager:

    ```
    # shutdown -t now “now”
    Broadcast message from root (pts/4) (Mon Mar 12 17:10:13 2012):
    The system is going down to maintenance mode NOW!
    ```

    When the command finishes, the server is powered off (standby).

Powering off the hardware

Power off the file serving nodes in any order. The step completely shuts down the cluster.

Starting the system

To start the system, first power on the file serving nodes, and then start the X900 Software.

Starting the StoreAll software

To start the StoreAll software, complete the following steps:

1. Power on the node hosting the active Fusion Manager.
2. Power on the file serving nodes (*root segment = segment 1; power on owner first, if possible).
3. Monitor the nodes on the GUI and wait for them all to report `UP` in the output from the following command:

   ```
   ibrix_server -l
   ```
4. Mount file systems and verify their content. Run the following command on the file serving node hosting the active Fusion Manager:

   ```
   ibrix_mount -f fs_name -m <mountpoint>
   ```

   On Linux StoreAll clients, run the following command:

   ```
   ibrix_lwmount -f fsname -m <mountpoint>
   ```
5. Enable HA on the file serving nodes. Run the following command on the file serving node hosting the active Fusion Manager:

   ```
   ibrix_server -m
   ```
6. From the active Fusion Manager, enter the following command to move all Fusion Managers into passive mode:

   ```
   ibrix_fm -m passive -A
   ```

The StoreAll software is now available, and you can now access your file systems.

Powering file serving nodes on or off

When file serving nodes are connected to properly configured power sources, the nodes can be powered on or off or can be reset remotely. To prevent interruption of service, set up standbys for the nodes (see "Configuring High Availability on the cluster" (page 54), and then manually fail them over before powering them off (see "Failing a server over manually" (page 64)). Remotely powering off a file serving node does not trigger failover.

To power on, power off, or reset a file serving node, use the following command:

```
ibrix_server -P {on|reset|off} -h HOSTNAME
```

Performing a rolling reboot

The rolling reboot procedure allows you to reboot all file serving nodes in the cluster while the cluster remains online. Before beginning the procedure, ensure that each file serving node has a backup node and that StoreAll HA is enabled. See "Configuring virtual interfaces for client access" (page 48) and "Configuring High Availability on the cluster" (page 54) for more information about creating standby backup pairs, where each server in a pair is the standby for the other.

Use one of the following schemes for the reboot:

- Reboot the file serving nodes one-at-a-time.
- Divide the file serving nodes into two groups, with the nodes in the first group having backups in the second group, and the nodes in the second group having backups in the first group. You can then reboot one group at-a-time.

To perform the rolling reboot, complete the following steps on each file serving node:

1. Reboot the node directly from Linux. (Do not use the "Power Off" functionality in the GUI, as it does not trigger failover of file serving services.) The node will fail over to its backup.
2. Wait for the GUI to report that the rebooted node is Up.
3. From the GUI, failback the node, returning services to the node from its backup. Run the following command on the backup node:

   ```
   ibrix_server -f -U -h HOSTNAME
   ```

   `HOSTNAME` is the name of the node that you just rebooted.

Starting and stopping processes

You can start, stop, and restart processes and can display status for the processes that perform internal StoreAll software functions. The following commands also control the operation of PostgreSQL on the machine. The PostgreSQL service is available at `/usr/local/ibrix/init/`.

To start and stop processes and view process status on the Fusion Manager, use the following command:

```
/etc/init.d/ibrix_fusionmanager [start | stop | restart | status]
```

To start and stop processes and view process status on a file serving node, use the following command. In certain situations, a follow-up action is required after stopping, starting, or restarting a file serving node.

```
/etc/init.d/ibrix_server [start | stop | restart | status]
```

To start and stop processes and view process status on a StoreAll client, use the following command:

```
/etc/init.d/ibrix_client [start | stop | restart | status]
```

Tuning file serving nodes and StoreAll clients

Typically, HP Support sets the tuning parameters on the file serving nodes during the cluster installation and changes should be needed only for special situations.

△ **CAUTION:** The default values for the host tuning parameters are suitable for most cluster environments. Because changing parameter values can alter file-system performance, HP recommends that you exercise caution before implementing any changes, or do so only under the guidance of HP technical support.

Host tuning changes are executed immediately for file serving nodes. For StoreAll clients, a tuning intention is stored in the Fusion Manager. When StoreAll software services start on a client, the client queries the Fusion Manager for the host tunings that it should use and then implements them. If StoreAll software services are already running on a client, you can force the client to query the Fusion Manager by executing `ibrix_client` or `ibrix_lwhost --a` on the client, or by rebooting the client.

You can locally override host tunings that have been set on StoreAll Linux clients by executing the `ibrix_lwhost` command.

Tuning file serving nodes on the GUI

The Modify Server(s) Wizard can be used to tune one or more servers in the cluster. To open the wizard, select **Servers** from the Navigator and then click **Tuning Options** from the Summary panel.

The General Tunings dialog box specifies the communications protocol (TCP or UDP) and the number of admin and server threads.

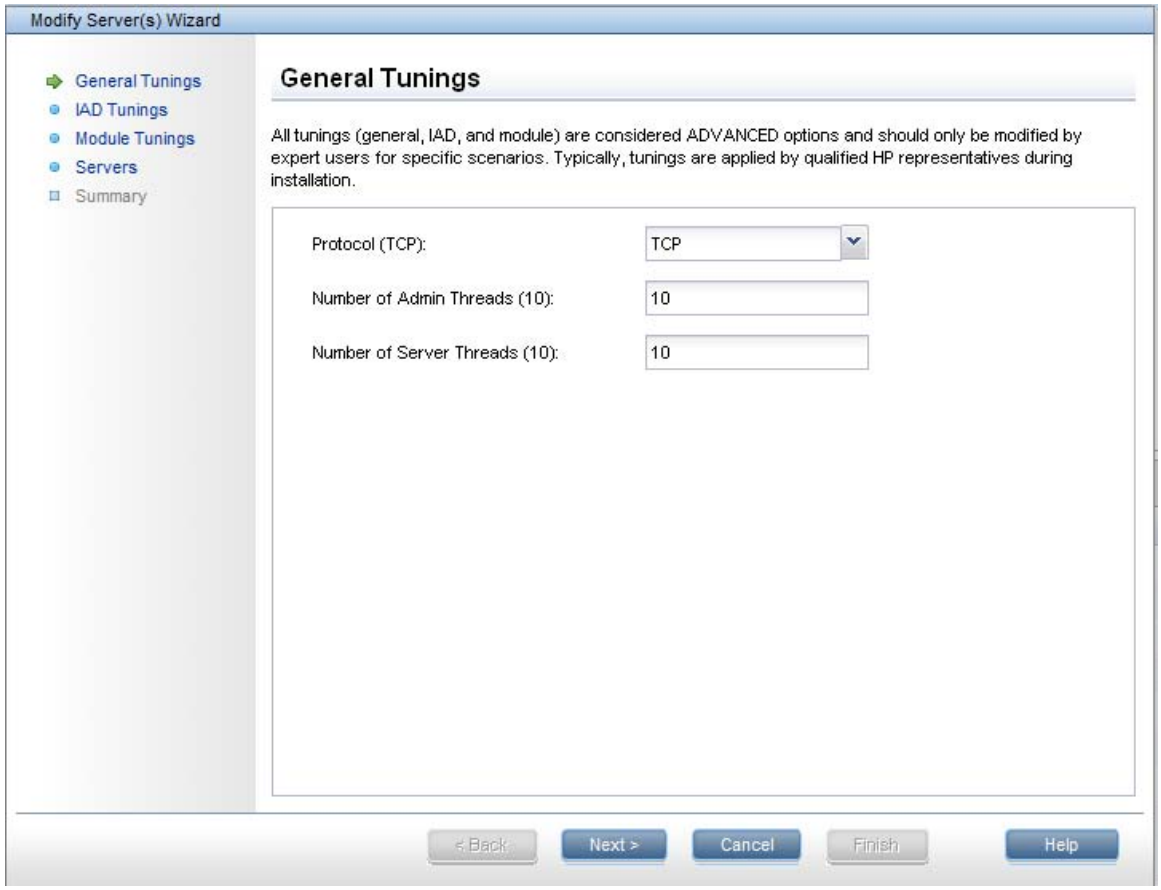

The IAD Tunings dialog box configures the StoreAll administrative daemon.

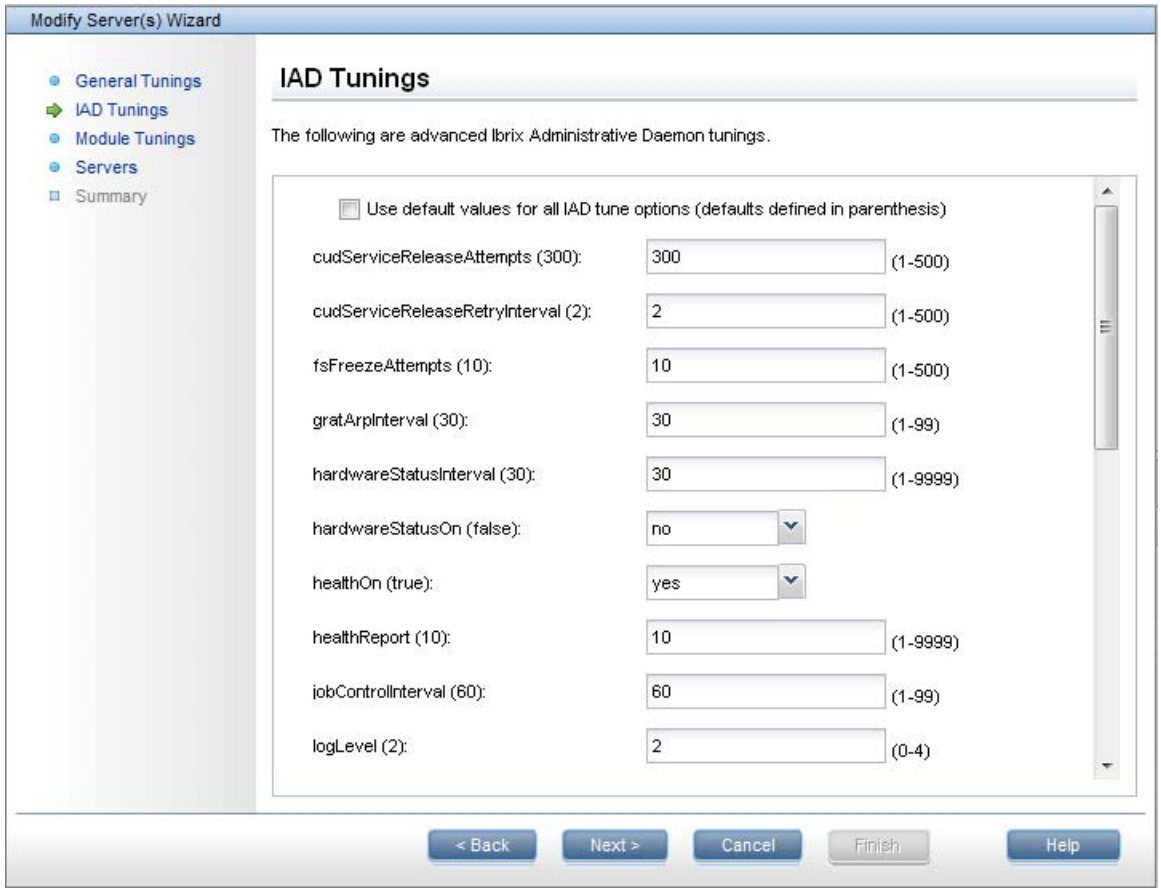

The Module Tunings dialog box adjusts various advanced parameters that affect server operations.

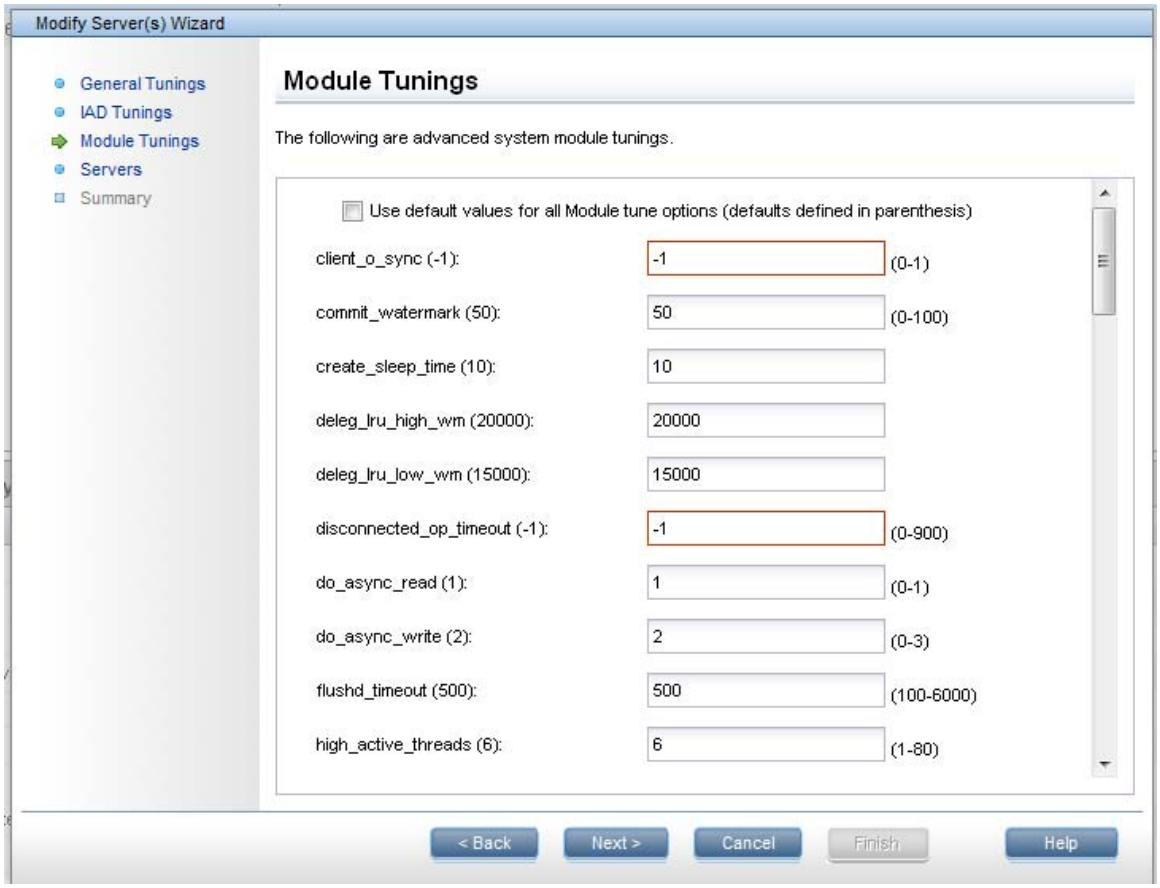

On the Servers dialog box, select the servers to which the tunings should be applied.

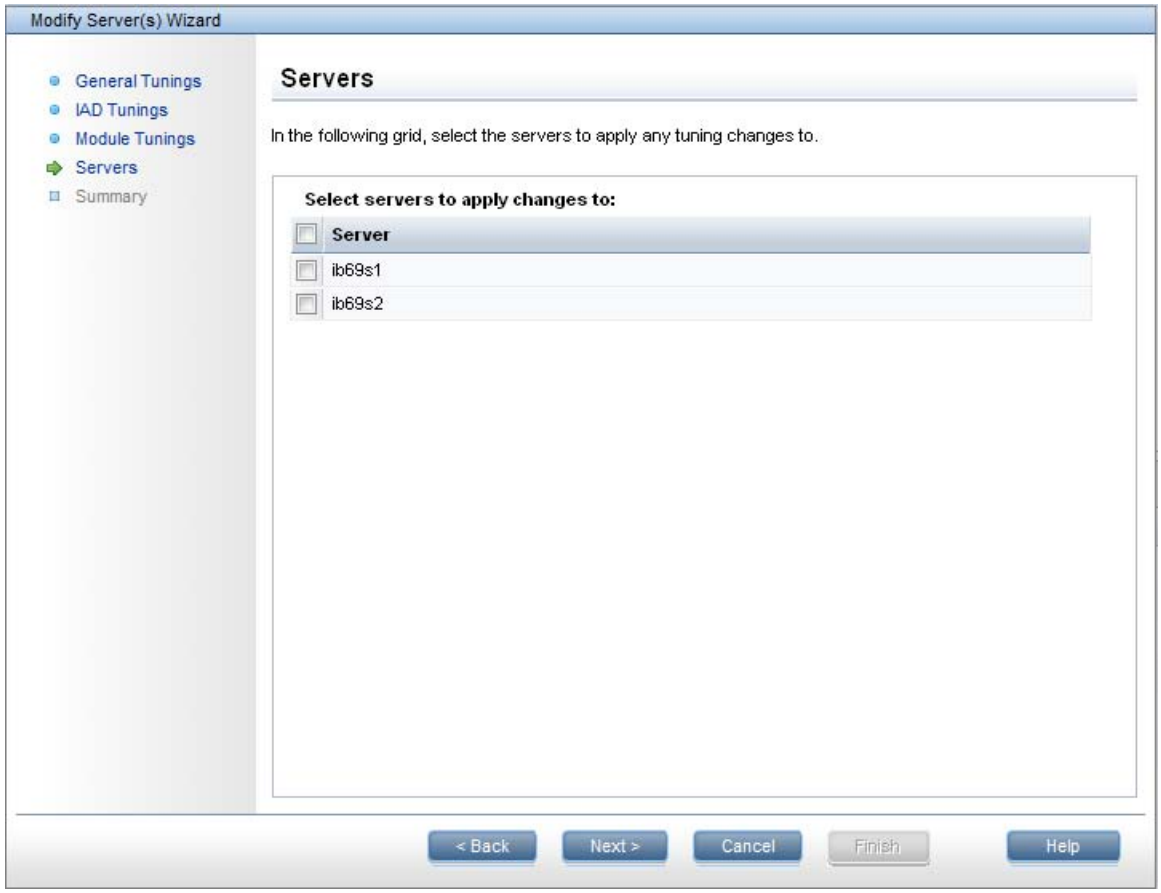

Tuning file serving nodes from the CLI

All Fusion Manager commands for tuning hosts include the `-h HOSTLIST` option, which supplies one or more host groups. Setting host tunings on a host group is a convenient way to tune a set of clients all at once. To set the same host tunings on all clients, specify the `clients` host group.

CAUTION: Changing host tuning settings alters file system performance. Contact HP Support before changing host tuning settings.

Use the `ibrix_host_tune` command to list or change host tuning settings:

- To list default values and valid ranges for all permitted host tunings:

  ```
  ibrix_host_tune -L
  ```

- To tune host parameters on nodes or host groups:

  ```
  ibrix_host_tune -S {-h HOSTLIST|-g GROUPLIST} -o OPTIONLIST
  ```

  Contact HP Support to obtain the values for *OPTIONLIST*. List the options as *option=value* pairs, separated by commas. To set host tunings on all clients, include the `-g clients` option.

- To reset host parameters to their default values on nodes or host groups:

  ```
  ibrix_host_tune -U {-h HOSTLIST|-g GROUPLIST} [-n OPTIONS]
  ```

  To reset all options on all file serving nodes, host groups, and StoreAll clients, omit the `-h HOSTLIST` and `-n OPTIONS` options. To reset host tunings on all clients, include the `-g clients` option.

The values that are restored depend on the values specified for the `-h HOSTLIST` command:

- **File serving nodes.** The default file serving node host tunings are restored.
- **StoreAll clients.** The host tunings that are in effect for the default `clients` host group are restored.
- **Hostgroups.** The host tunings that are in effect for the parent of the specified host groups are restored.

- To list host tuning settings on file serving nodes, StoreAll clients, and host groups, use the following command. Omit the `-h` argument to see tunings for all hosts. Omit the `-n` argument to see all tunings.

  ```
  ibrix_host_tune -l [-h HOSTLIST] [-n OPTIONS]
  ```

- To set the communications protocol on nodes and host groups, use the following command. To set the protocol on all StoreAll clients, include the `-g clients` option.
- `ibrix_host_tune -p {UDP|TCP} {-h HOSTLIST| -g GROUPLIST}`
- To set server threads on file serving nodes, host groups, and StoreAll clients:

  ```
  ibrix_host_tune -t THREADCOUNT {-h HOSTLIST| -g GROUPLIST}
  ```

- To set admin threads on file serving nodes, host groups, and StoreAll clients, use this command. To set admin threads on all StoreAll clients, include the `-g clients` option.

  ```
  ibrix_host_tune -a THREADCOUNT {-h HOSTLIST| -g GROUPLIST}
  ```

Tuning StoreAll clients locally

Linux clients. Use the `ibrix_lwhost` command to tune host parameters. For example, to set the communications protocol:

```
ibrix_lwhost --protocol -p {tcp|udp}
```

To list host tuning parameters that have been changed from their defaults:

```
ibrix_lwhost --list
```

See the `ibrix_lwhost` command description in the *HP StoreAll Storage CLI Reference Guide* for other available options.

Windows clients. Click the **Tune Host** tab on the Windows StoreAll client GUI. Tunable parameters include the NIC to prefer (the default is the cluster interface), the communications protocol (UDP or TCP), and the number of server threads to use. See the online help for the client if necessary.

Managing segments

When a file system is created, the servers accessing the file system are assigned ownership of the storage segments used for the file system. Each server is responsible for managing the segments it owns.

When the cluster is expanded, the StoreAll software attempts to maintain proper load balancing and utilization in the following ways:

- When servers are added, ownership of the existing segments is redistributed among the available servers.
- When storage is added, ownership of the new segments is distributed among the available servers.

Occasionally you may need to manage the segments manually:

- **Migrate segments.** This operation transfers ownership of segments to other servers. For example, if a server is overloaded or unavailable, you can transfer its segments to another server that can see the same storage.
- **Rebalance segments.** This operation redistributes files across segments and can be used if certain segments are filling up and affecting file system performance. See "Maintaining file systems" in the for more information.
- **Evacuate segments.** This operation moves the data in a segment to another segment. It is typically used before removing storage from the cluster.

Migrating segments

Segment migration transfers segment ownership but it does not move segments from their physical locations in the storage system. Segment ownership is recorded on the physical segment itself, and the ownership data is part of the metadata that the Fusion Manager distributes to file serving nodes and StoreAll clients so that they can locate segments.

To migrate segments on the GUI, select the file system on the Filesystems panel, select **Segments** from the lower Navigator, and then click **Ownership/Migration** on the Segments panel to open the Segment Ownership Migration Wizard.

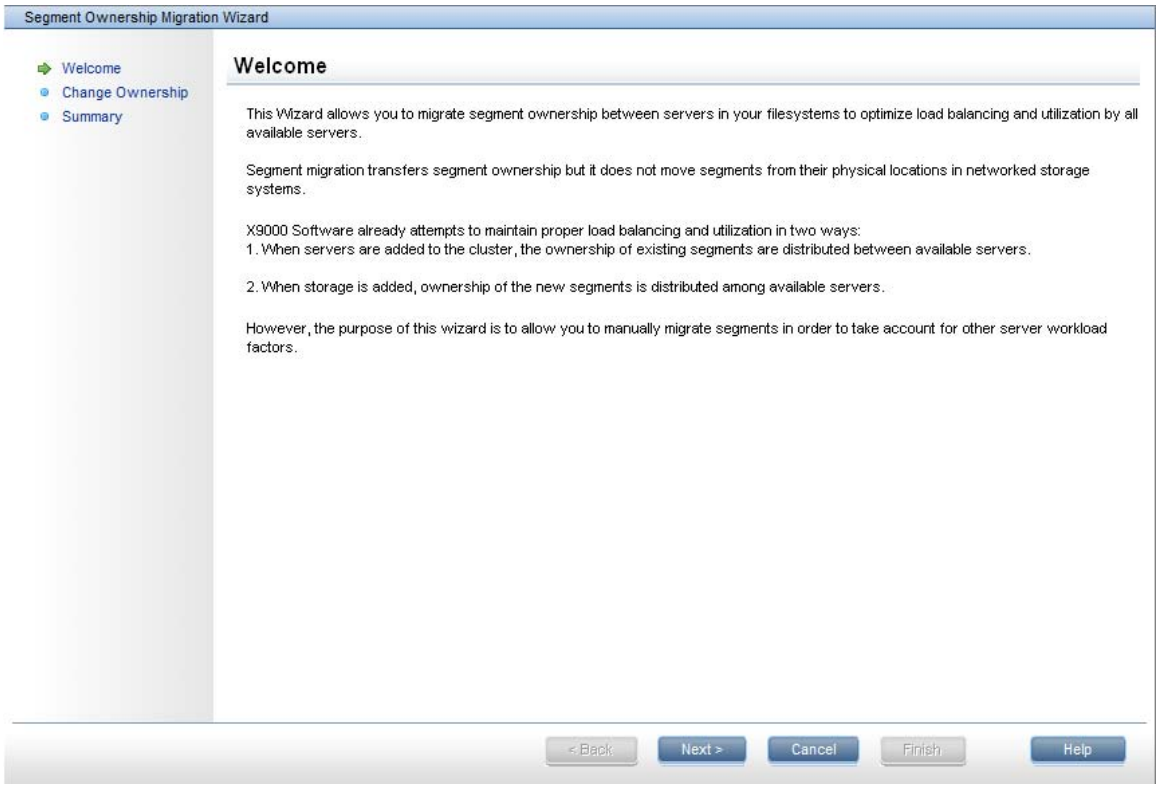
Segment Ownership Migration Wizard
Welcome
Change Ownership
Summary
Welcome
This Wizard allows you to migrate segment ownership between servers in your filesystems to optimize load balancing and utilization by all available servers.
Segment migration transfers segment ownership but it does not move segments from their physical locations in networked storage systems.
X9000 Software already attempts to maintain proper load balancing and utilization in two ways:
1. When servers are added to the cluster, the ownership of existing segments are distributed between available servers.
2. When storage is added, ownership of the new segments is distributed among available servers.
However, the purpose of this wizard is to allow you to manually migrate segments in order to take account for other server workload factors.
< Back
Next >
Cancel
Finish
Help

The Change Ownership dialog box reports the status of the servers in the cluster and lists the segments owned by each server. In the Segment Properties section of the dialog box, select the segment whose ownership you are transferring, and click **Change Owner**.

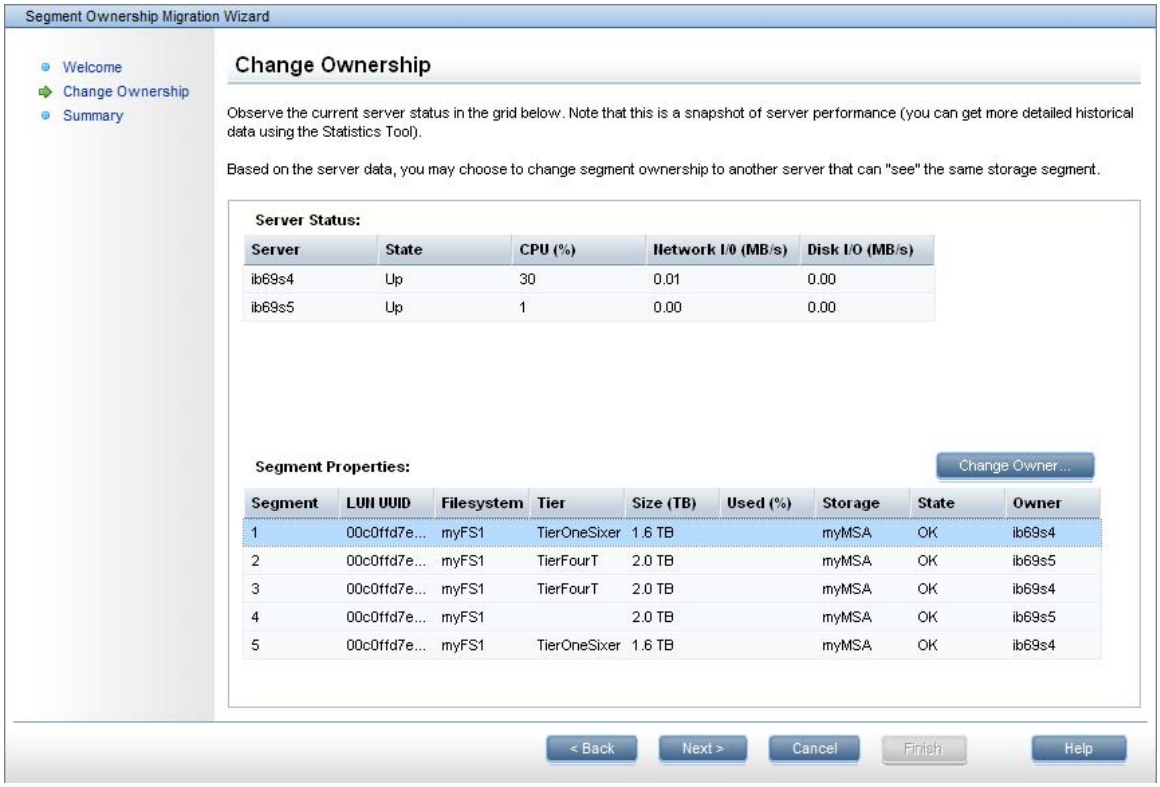
Segment Ownership Migration Wizard
Welcome
Change Ownership
Summary
Change Ownership
Observe the current server status in the grid below. Note that this is a snapshot of server performance (you can get more detailed historical data using the Statistics Tool).
Based on the server data, you may choose to change segment ownership to another server that can "see" the same storage segment.
Server Status:
Server State CPU (%) Network I/0 (MB/s) Disk I/O (MB/s)
ib69s4 Up 30 0.01 0.00
ib69s5 Up 1 0.00 0.00
Segment Properties:
Change Owner...
Segment LUN UUID Filesystem Tier Size (TB) Used (%) Storage State Owner
1 00c0ffd7e... myFS1 TierOneSixer 1.6 TB myMSA OK ib69s4
2 00c0ffd7e... myFS1 TierFourT 2.0 TB myMSA OK ib69s5
3 00c0ffd7e... myFS1 TierFourT 2.0 TB myMSA OK ib69s4
4 00c0ffd7e... myFS1 2.0 TB myMSA OK ib69s5
5 00c0ffd7e... myFS1 TierOneSixer 1.6 TB myMSA OK ib69s4
< Back
Next >
Cancel
Finish
Help

The new owner of the segment must be able to see the same storage as the original owner. The Change Segment Owner dialog box lists the servers that can see the segment you selected. Select one of these servers to be the new owner.

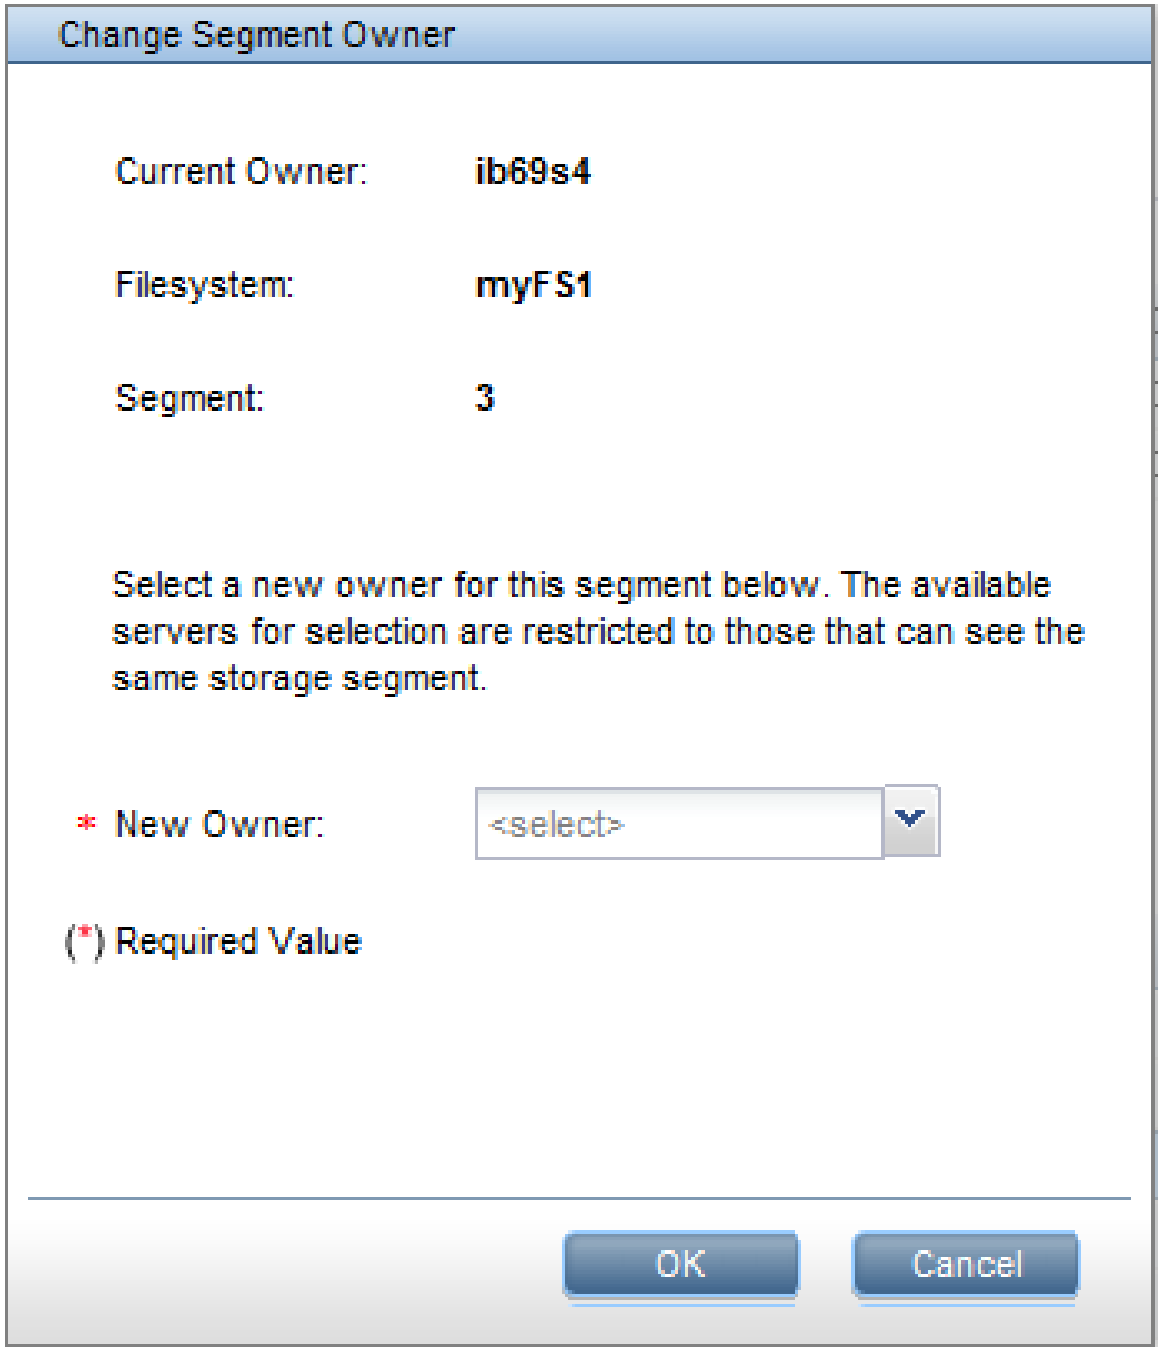

The Summary dialog box shows the segment migration you specified. Click **Back** to make any changes, or click **Finish** to complete the operation.

To migrate ownership of segments from the CLI, use the following commands.

Migrate ownership of specific segments:

```
ibrix_fs -m -f FSNAME -s LVLIST -h HOSTNAME [-M] [-F] [-N]
```

To force the migration, include `-M`. To skip the source host update during the migration, include `-F`. To skip server health checks, include `-N`.

The following command migrates ownership of segments `ilv2` and `ilv3` in file system `ifs1` to `server2`:

```
ibrix_fs -m -f ifs1 -s ilv2,ilv3 -h server2
```

Migrate ownership of all segments owned by specific servers:

```
ibrix_fs -m -f FSNAME -H HOSTNAME1,HOSTNAME2 [-M] [-F] [-N]
```

For example, to migrate ownership of all segments in file system `ifs1` from `server1` to `server2`:

```
ibrix_fs -m -f ifs1 -H server1,server2
```

Evacuating segments and removing storage from the cluster

Before removing storage used for a StoreAll software file system, you will need to evacuate the segments (or logical volumes) storing file system data. This procedure moves the data to other segments in the file system and is transparent to users or applications accessing the file system. When evacuating a segment, you should be aware of the following restrictions:

- While the evacuation task is running, the system prevents other tasks from running on the file system. Similarly, if another task is running on the file system, the evacuation task cannot be scheduled until the first task is complete.
- The file system must be quiescent (no active I/O while a segment is being evacuated). Running this utility while the file system is active may result in data inconsistency or loss.

To evacuate a segment, complete the following steps:

1. Identify the segment residing on the physical volume to be removed. Select **Storage** from the Navigator on the GUI. Note the file system and segment number on the affected physical volume.
2. Locate other segments on the file system that can accommodate the data being evacuated from the affected segment. Select the file system on the GUI and then select **Segments** from the lower Navigator. If segments with adequate space are not available, add segments to the file system.
3. Evacuate the segment. Select the file system on the GUI, select **Segments** from the lower Navigator, and then click **Rebalance/Evacuate** on the Segments panel. When the Segment Rebalance and Evacuation Wizard opens, select **Evacuate** as the mode.

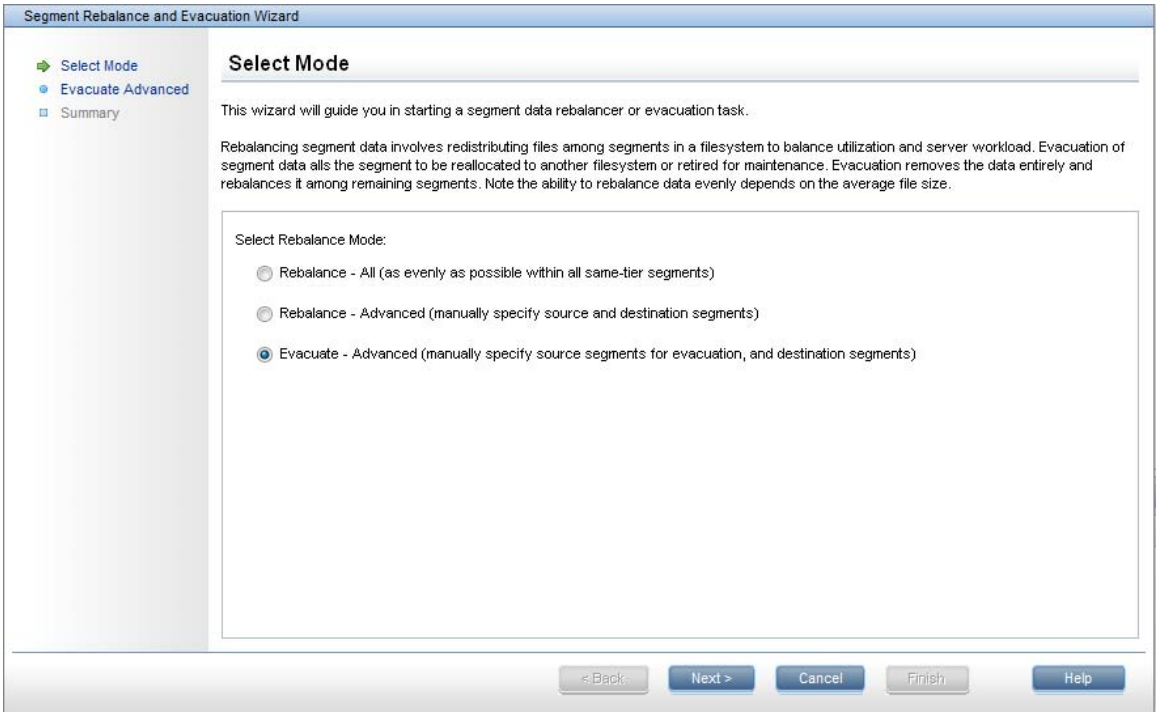
Segment Rebalance and Evacuation Wizard
Select Mode
Evacuate Advanced
Summary
Select Mode
This wizard will guide you in starting a segment data rebalancer or evacuation task.
Rebalancing segment data involves redistributing files among segments in a filesystem to balance utilization and server workload. Evacuation of segment data alls the segment to be reallocated to another filesystem or retired for maintenance. Evacuation removes the data entirely and rebalances it among remaining segments. Note the ability to rebalance data evenly depends on the average file size.
Select Rebalance Mode:
Rebalance - All (as evenly as possible within all same-tier segments)
Rebalance - Advanced (manually specify source and destination segments)
Evacuate - Advanced (manually specify source segments for evacuation, and destination segments)
< Back
Next >
Cancel
Finish
Help

On the Evacuate Advanced dialog box, locate the segment to be evacuated and click **Source**. Then locate the segments that will receive the data from the segment and click **Destination**. If the file system is tiered, be sure to select destination segments on the same tier as the source segment.

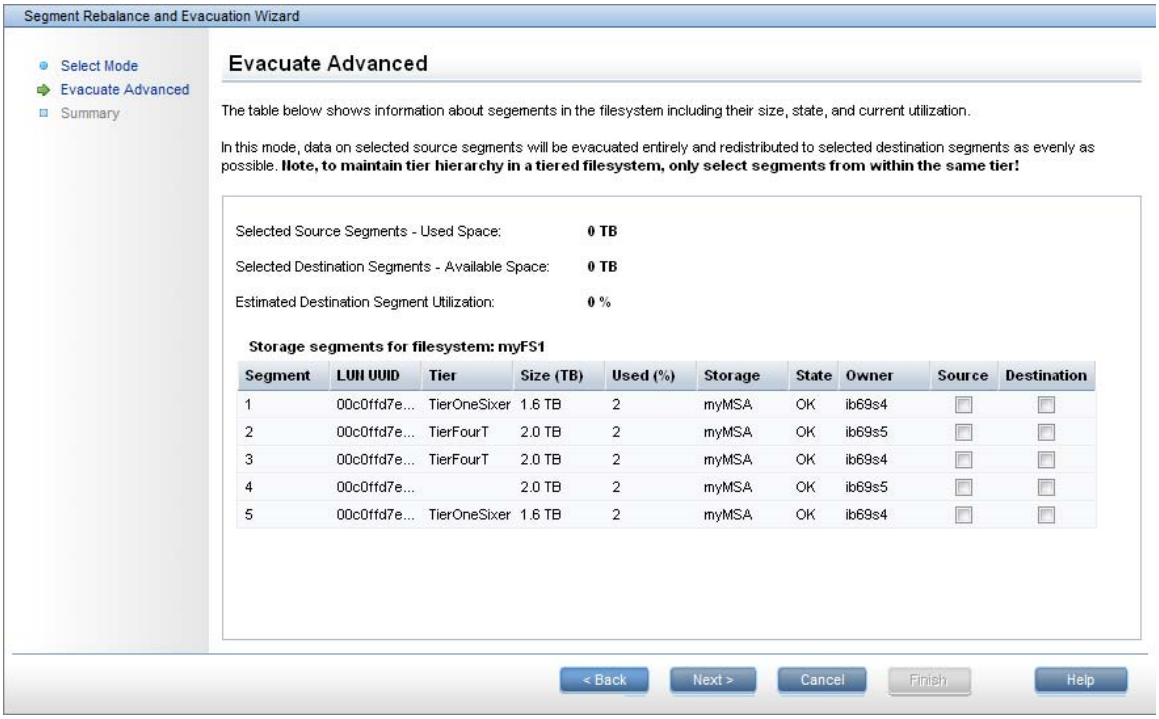
Segment Rebalance and Evacuation Wizard
Select Mode
Evacuate Advanced
Summary
Evacuate Advanced
The table below shows information about segements in the filesystem including their size, state, and current utilization.
In this mode, data on selected source segments will be evacuated entirely and redistributed to selected destination segments as evenly as possible. Note, to maintain tier hierarchy in a tiered filesystem, only select segments from within the same tier!
Selected Source Segments - Used Space: 0 TB
Selected Destination Segments - Available Space: 0 TB
Estimated Destination Segment Utilization: 0 %
Storage segments for filesystem: myFS1
Segment LUN UUID Tier Size (TB) Used (%) Storage State Owner Source Destination
1 00c0ffd7e... TierOneSixer 1.6 TB 2 myMSA OK ib69s4
2 00c0ffd7e... TierFourT 2.0 TB 2 myMSA OK ib69s5
3 00c0ffd7e... TierFourT 2.0 TB 2 myMSA OK ib69s4
4 00c0ffd7e... 2.0 TB 2 myMSA OK ib69s5
5 00c0ffd7e... TierOneSixer 1.6 TB 2 myMSA OK ib69s4
< Back
Next >
Cancel
Finish
Help

The Summary dialog box lists the source and destination segments for the evacuation. Click **Back** to make any changes, or click **Finish** to start the evacuation.

The Active Tasks panel reports the status of the evacuation task. When the task is complete, it will be added to the Inactive Tasks panel.

4. When the evacuation is complete, run the following command to retire the segment from the file system:

   ```
   ibrix_fs -B -f FSNAME -n BADSEGNUMLIST
   ```

   The segment number associated with the storage is not reused. The underlying LUN or volume can be reused in another file system or physically removed from the storage solution when this step is complete.

5. If quotas were disabled on the file system, unmount the file system and then re-enable quotas using the following command:

   ```
   ibrix_fs -q -E -f FSNAME
   ```

   Then remount the file system.

To evacuate a segment using the CLI, use the `ibrix_evacuate` command, as described in the *HP StoreAll Storage CLI Reference Guide*.

Troubleshooting segment evacuation

- If segment evacuation fails, HP recommends that you run phase 1 of the `ibrix_fsck` command in corrective mode on the segment that failed the evacuation. For more information, see "Checking and repairing file systems" in the *HP StoreAll Storage File System User Guide*.
- The segment evacuation process fails if a segment contains chunk files bigger than 3.64 T; you need to move these chunk files manually. The evacuation process generates a log reporting the chunk files on the segment that were not moved. The log file is saved in the management console log directory (the default is `/usr/local/ibrix/log`) and is named `Rebalance_<jobID>-<FS-ID>.info` (for example, `Rebalance_29-ibfs1.info`).

  Run the `inum2name` command to identify the symbolic name of the chunk file:

  ```
  # ./inum2name --fsname=ibfs 500000017
  ibfs:/sliced_dir/file3.bin
  ```

After obtaining the name of the file, use a command such as `cp` to move the file manually. Then run the segment evacuation process again.

The analyzer log lists the chunks that were left on segments. Following is an example of the log:

```
2012-03-13 11:57:35:0332834 | <INFO> | 1090169152 | segment 3 not migrated
chunks 462
2012-03-13 11:57:35:0332855 | <INFO> | 1090169152 | segment 3 not migrated
replicas 0
2012-03-13 11:57:35:0332864 | <INFO> | 1090169152 | segment 3 not migrated
files 0
2012-03-13 11:57:35:0332870 | <INFO> | 1090169152 | segment 3 not migrated
directories 0
2012-03-13 11:57:35:0332875 | <INFO> | 1090169152 | segment 3 not migrated root
0
2012-03-13 11:57:35:0332880 | <INFO> | 1090169152 | segment 3 orphan inodes 0
2012-03-13 11:57:35:0332886 | <INFO> | 1090169152 | segment 3 chunk: inode
3099CC002.8E2124C4, poid 3099CC002.8E2124C4, primary 807F5C010.36B5072B poid
807F5C010.36B5072B
2012-03-13 11:57:35:0332894 | <INFO> | 1090169152 | segment 3 chunk: inode
3099AC007.8E2125A1, poid 3099AC007.8E2125A1, primary 60A1D8024.42966361 poid
60A1D8024.42966361
2012-03-13 11:57:35:0332901 | <INFO> | 1090169152 | segment 3 chunk: inode
3015A4031.C34A99FA, poid 3015A4031.C34A99FA, primary 40830415E.7793564B poid
40830415E.7793564B
2012-03-13 11:57:35:0332908 | <INFO> | 1090169152 | segment 3 chunk: inode
3015A401B.C34A97F8, poid 3015A401B.C34A97F8, primary 4083040D9.77935458 poid
4083040D9.77935458
2012-03-13 11:57:35:0332915 | <INFO> | 1090169152 | segment 3 chunk: inode
3015A4021.C34A994C, poid 3015A4021.C34A994C, primary 4083040FF.7793558E poid
4083040FF.7793558E
```

Use the `inum2name` utility to translate the primary inode ID into the file name.

Removing a node from a cluster

In the following procedure, the cluster contains four nodes: `FSN1`, `FSN2`, `FSN3`, and `FSN4`. `FSN4` is the node being removed. The user NIC for `FSN4` is `bond0:1`. The file system name is `ibfs1`, which is mounted on `/ibfs1` and shared as `ibfs1` through NFS and SMB . `FSN3` and `FSN4` are the failover pair, and `bond0:2` is configured as the stand-by interface.

1. Stop High Availability:

   ```
   ibrix_server -m –U
   ```

2. Verify that the Active Fusion Manager is on a server other than `FSN4`. Run the following command from `FSN4`:

   ```
   ibrix_fm -i
   ```

   If the Active Fusion Manager is on FSN4, move the Fusion Manager to nofmfailover mode:

   ```
   ibrix_fm -m nofmfailover
   ```

3. Migrate all segments from `FSN4` to its failover partner (`FSN3`):

   ```
   ibrix_fs -m -f ibfs1 -h FSN4,FSN3
   ```

   NOTE: If there is a large number of segments to be migrated and/or segments can be migrated to several FSNs (in SAN environments), run the following command instead:

   ```
   ibrix_fs -m -f ibfs1 -s LVLIST -h FSN_name
   ```

4. Remove the High Availability configuration between FSN4 and FSN3:
   a. Stop NIC monitoring for the user NICs:
      ```
      ibrix_nic -m -h FSN3 -D FSN4/bond0:1
      ibrix_nic -m -h FSN4 -D FSN3/bond0:1
      ```
   b. Remove the backup NICs:
      ```
      ibrix_nic -b -u FSN4/bond0:2
      ibrix_nic -b -u FSN3/bond0:2
      ```
   c. Remove the backup server:
      ```
      ibrix_server -b -U -h FSN4
      ibrix_server -b -U -h FSN3
      ```
5. If FSN4 is configured for DNS round robin or there are NFS/SMB clients mounting file systems directly from FSN4, you must migrate the user NIC IP address to another FSN. Continue with step 6 if the FSN4 NIC IP address has been removed and there are no clients accessing data from FSN4.
   a. Create a new placeholder NIC on FSN3 you will be migrating IP address to:
      ```
      ibrix_nic -a -n bond0:3 -h FSN3
      ```
   b. Migrate NIC bond0:1 from FSN4 to bond0:3 on FSN3:
      ```
      ibrix_nic -s -H bond0:1/FSN4,bond0:3/FSN3
      ```
6. Stop NFS and SMB services on FSN4:
   ```
   ibrix_server -t -s NFS -c stop
   ibrix_server –t –s cifs –c stop
   ```
7. Remove all NFS and SMB shares from FSN4 (in this example, ibfs1 is shared via NFS and CIFS):
   ```
   ibrix_exportfs -U -h FSN4 -p *:/ibfs1
   ibrix_cifs -d -s ibfs1 -h FSN4
   ```
8. Unmount ibfs1 from FSN4 and delete the mountpoint on FSN4 from the cluster:
   ```
   ibrix_umount –f ibfs1 –h FSN4
   ibrix_mountpoint –d –h FSN4 –m /ibfs1
   ```
9. Remove FSN4 from AgileFM quorum participation:
   ```
   ibrix_fm -u FSN4
   ```
10. Delete FSN4 from the cluster:
    ```
    ibrix_server -d -h FSN4
    ```
11. Reconfigure High Availability on FSN3, if needed.
12. Enable HA:
    ```
    ibrix_server -m
    ```
13. Save the new cluster configuration:
    ```
    ibrix_fm -B
    ```

14. Uninstall the StoreALL OS software from `FSN4`:

    ```
    /usr/local/ibrix/local/installation/ibrix/ibrixinit -u -F
    ```

    NOTE: If the same StoreAll OS version will be reinstalled on `FSN4`, use the following command instead:

    ```
    /usr/local/ibrix/local/installation/ibrix/ibrixinit -U
    ```

The node is no longer in the cluster.

Maintaining networks

Cluster and user network interfaces

StoreAll software supports the following logical network interfaces:

- **Cluster network interface.** This network interface carries Fusion Manager traffic, traffic between file serving nodes, and traffic between file serving nodes and clients. A cluster can have only one cluster interface. For backup purposes, each file serving node can have two cluster NICs.
- **User network interface.** This network interface carries traffic between file serving nodes and clients. Multiple user network interfaces are permitted.

The cluster network interface was created for you when your cluster was installed. (A virtual interface is used for the cluster network interface.) One or more user network interfaces may also have been created, depending on your site's requirements. You can add user network interfaces as necessary.

Adding user network interfaces

Although the cluster network can carry traffic between file serving nodes and either NFS/SMB/HTTP/FTP or StoreAll clients, you may want to create user network interfaces to carry this traffic. If your cluster must accommodate a mix of NFS/SMB/FTP/HTTP clients and StoreAll clients, or if you need to segregate client traffic to different networks, you will need one or more user networks. In general, it is better to assign a user network for protocol (NFS/SMB/HTTP/FTP) traffic because the cluster network cannot host the virtual interfaces (VIFs) required for failover. HP recommends that you use a Gigabit Ethernet port (or faster) for user networks.

When creating user network interfaces for file serving nodes, keep in mind that nodes needing to communicate for file system coverage or for failover must be on the same network interface. Also, nodes set up as a failover pair must be connected to the same network interface.

For a highly available cluster, HP recommends that you put protocol traffic on a user network and then set up automated failover for it (see "Configuring High Availability on the cluster" (page 54)). This method prevents interruptions to the traffic. If the cluster interface is used for protocol traffic and that interface fails on a file serving node, any protocol clients using the failed interface to access a mounted file system will lose contact with the file system because they have no knowledge of the cluster and cannot reroute requests to the standby for the node.

Link aggregation and virtual interfaces

When creating a user network interface, you can use link aggregation to combine physical resources into a single VIF. VIFs allow you to provide many named paths within the larger physical resource, each of which can be managed and routed independently, as shown in the following diagram. See the network interface vendor documentation for any rules or restrictions required for link aggregation.

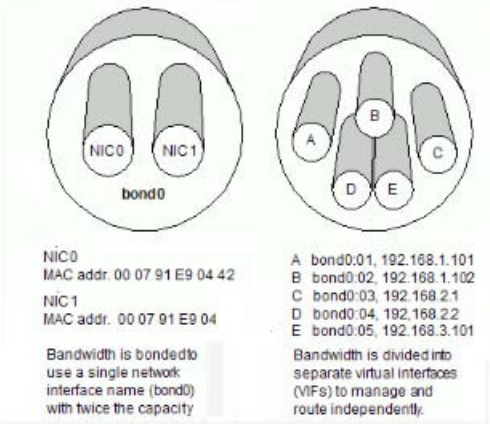

Identifying a user network interface for a file serving node

To identify a user network interface for specific file serving nodes, use the `ibrix_nic` command. The interface name (*IFNAME*) can include only alphanumeric characters and underscores, such as `eth1`.

```
ibrix_nic -a -n IFNAME -h HOSTLIST
```

If you are identifying a VIF, add the VIF suffix (*:nnnn*) to the physical interface name. For example, the following command identifies virtual interface `eth1:1` to physical network interface `eth1` on file serving nodes `s1.hp.com` and `s2.hp.com`:

```
ibrix_nic -a -n eth1:1 -h s1.hp.com,s2.hp.com
```

When you identify a user network interface for a file serving node, the Fusion Manager queries the node for its IP address, netmask, and MAC address and imports the values into the configuration database. You can modify these values later if necessary.

If you identify a VIF, the Fusion Manager does not automatically query the node. If the VIF will be used only as a standby network interface in an automated failover setup, the Fusion Manager will query the node the first time a network is failed over to the VIF. Otherwise, you must enter the VIF's IP address and netmask manually in the configuration database (see "Setting network interface options in the configuration database" (page 123)). The Fusion Manager does not require a MAC address for a VIF.

If you created a user network interface for StoreAll client traffic, you will need to prefer the network for the StoreAll clients that will use the network (see "Preferring network interfaces" (page 123)).

Setting network interface options in the configuration database

To make a VIF usable, execute the following command to specify the IP address and netmask for the VIF. You can also use this command to modify certain `ifconfig` options for a network.

```
ibrix_nic -c -n IFNAME -h HOSTNAME [-I IPADDR] [-M NETMASK]
[-B BCASTADDR] [-T MTU]
```

For example, to set netmask 255.255.0.0 and broadcast address 10.0.0.4 for interface eth3 on file serving node `s4.hp.com`:

```
ibrix_nic -c -n eth3 -h s4.hp.com -M 255.255.0.0 -B 10.0.0.4
```

Preferring network interfaces

After creating a user network interface for file serving nodes or StoreAll clients, you will need to *prefer* the interface for those nodes and clients. (It is not necessary to prefer a network interface for NFS or SMB clients, because they can select the correct user network interface at mount time.)

A network interface preference is executed immediately on file serving nodes. For StoreAll clients, the preference intention is stored on the Fusion Manager. When StoreAll software services start on a client, the client queries the Fusion Manager for the network interface that has been preferred for it and then begins to use that interface. If the services are already running on StoreAll clients

when you prefer a network interface, you can force clients to query the Fusion Manager by executing the command `ibrix_lwhost --a` on the client or by rebooting the client.

Preferring a network interface for a file serving node or Linux StoreAll client

The first command prefers a network interface for a File Server Node; the second command prefers a network interface for a client.

```
ibrix_server -n -h SRCHOST -A DESTHOST/IFNAME
ibrix_client -n -h SRCHOST -A DESTHOST/IFNAME
```

Execute this command once for each destination host that the file serving node or StoreAll client should contact using the specified network interface (*IFNAME*). For example, to prefer network interface `eth3` for traffic from file serving node s1.hp.com to file serving node `s2.hp.com`:

```
ibrix_server -n -h s1.hp.com -A s2.hp.com/eth3
```

Preferring a network interface for a Windows StoreAll client

If multiple user network interfaces are configured on the cluster, you will need to select the preferred interface for this client. On the Windows StoreAll client GUI, specify the interface on the Tune Host tab, as in the following example.

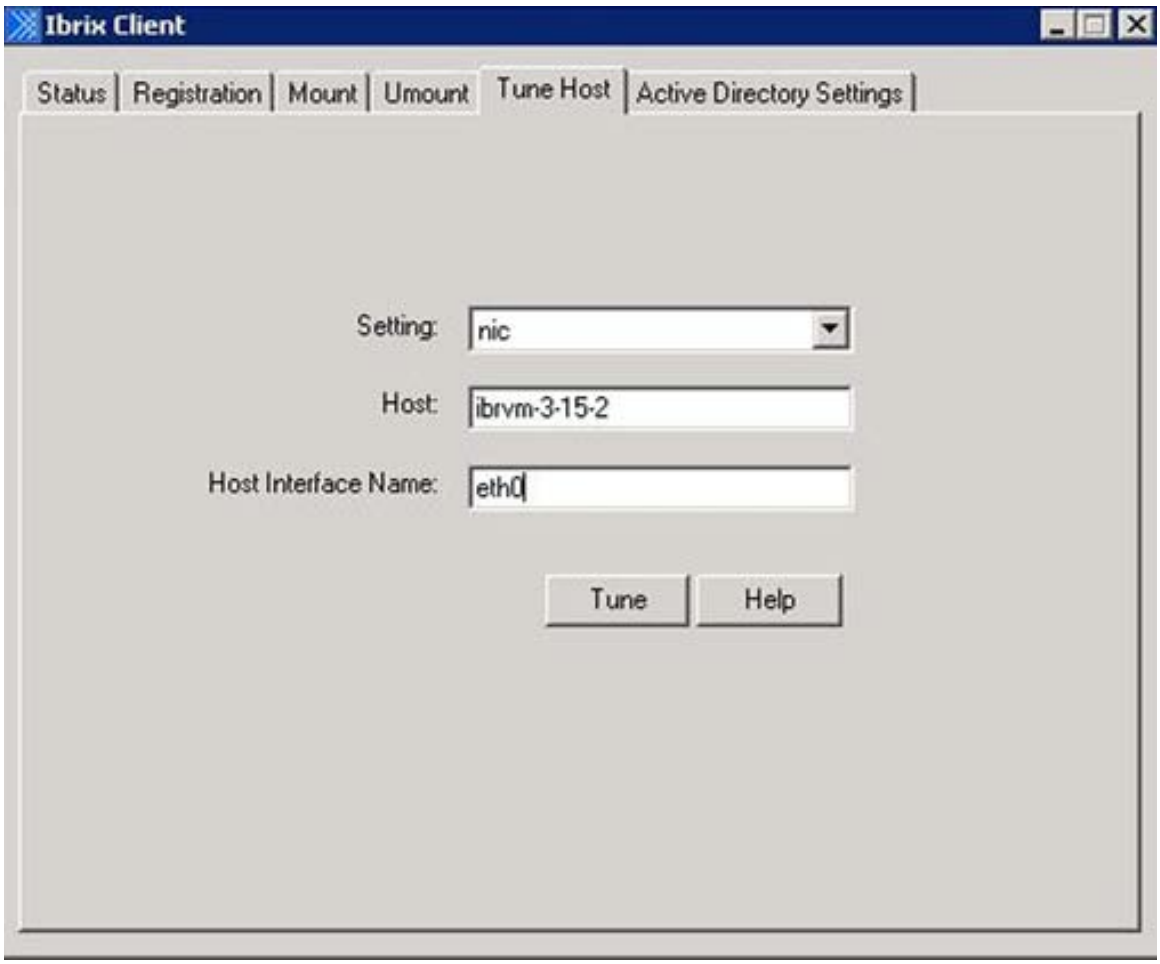

Preferring a network interface for a host group

You can prefer an interface for multiple StoreAll clients at one time by specifying a host group. To prefer a user network interface for all StoreAll clients, specify the `clients` host group. After preferring a network interface for a host group, you can locally override the preference on individual StoreAll clients with the command `ibrix_lwhost`.

To prefer a network interface for a host group, use the following command:

```
ibrix_hostgroup -n -g HOSTGROUP -A DESTHOST/IFNAME
```

The destination host (*DESTHOST*) cannot be a host group. For example, to prefer network interface `eth3` for traffic from all StoreAll clients (the `clients` host group) to file serving node `s2.hp.com`:

```
ibrix_hostgroup -n -g clients -A s2.hp.com/eth3
```

Unpreferring network interfaces

To return file serving nodes or StoreAll clients to the cluster interface, unprefer their preferred network interface. The first command unprefers a network interface for a file serving node; the second command unprefers a network interface for a client.

```
ibrix_server -n -h SRCHOST -D DESTHOST
ibrix_client -n -h SRCHOST -D DESTHOST
```

To unprefer a network interface for a host group, use the following command:

```
ibrix_client -n -g HOSTGROUP -A DESTHOST
```

Making network changes

This section describes how to change IP addresses, change the cluster interface, manage routing table entries, and delete a network interface.

Changing the IP address for a Linux StoreAll client

After changing the IP address for a Linux StoreAll client, you must update the StoreAll software configuration with the new information to ensure that the Fusion Manager can communicate with the client. Use the following procedure:

1. Unmount the file system from the client.
2. Change the client's IP address.
3. Reboot the client or restart the network interface card.
4. Delete the old IP address from the configuration database:

   ```
   ibrix_client -d -h CLIENT
   ```

5. Re-register the client with the Fusion Manager:

   ```
   register_client -p console_IPAddress -c clusterIF -n ClientName
   ```

6. Remount the file system on the client.

Changing the IP address for the cluster interface on a dedicated management console

You must change the IP address for the cluster interface on both the file serving nodes and the management console.

1. If High Availability is enabled, disable it by executing `ibrix_server -m -U`.
2. Unmount the file system from all file serving nodes, and reboot.
3. On each file serving node, locally change the IP address of the cluster interface.
4. Change the IP address of the cluster interface for each file serving node:

   ```
   <installdirectory>/bin/ibrix_nic -c -n IFNAME -h HOSTNAME [-I IPADDR]
   ```

5. Remount the file system.
6. Re-enable High Availability if necessary by executing `ibrix_server -m`.

Changing the cluster interface

If you restructure your networks, you might need to change the cluster interface. The following rules apply when selecting a new cluster interface:

- The Fusion Manager must be connected to all machines (including standby servers) that use the cluster network interface. Each file serving node and StoreAll client must be connected to

the Fusion Manager by the same cluster network interface. A Gigabit (or faster) Ethernet port must be used for the cluster interface.

- StoreAll clients must have network connectivity to the file serving nodes that manage their data and to the standbys for those servers. This traffic can use the cluster network interface or a user network interface.

To specify a new virtual cluster interface, use the following command:

```
ibrix_fm -c <VIF IP address> -d <VIF Device> -n <VIF Netmask>
-v cluster [-I <Local IP address_or_DNS hostname>]
```

Managing routing table entries

StoreAll Software supports one route for each network interface in the system routing table. Entering a new route for an interface overwrites the existing routing table entry for that interface.

Adding a routing table entry

To add a routing table entry, use the following command:

```
ibrix_nic -r -n IFNAME -h HOSTNAME -A -R ROUTE
```

The following command adds a route for virtual interface `eth2:232` on file serving node `s2.hp.com`, sending all traffic through gateway `gw.hp.com`:

```
ibrix_nic -r -n eth2:232 -h s2.hp.com -A -R gw.hp.com
```

Deleting a routing table entry

If you delete a routing table entry, it is not replaced with a default entry. A new replacement route must be added manually. To delete a route, use the following command:

```
ibrix_nic -r -n IFNAME -h HOSTNAME -D
```

The following command deletes all routing table entries for virtual interface `eth0:1` on file serving node `s2.hp.com`:

```
ibrix_nic -r -n eth0:1 -h s2.hp.com -D
```

Deleting a network interface

Before deleting the interface used as the cluster interface on a file serving node, you must assign a new interface as the cluster interface. See "Changing the cluster interface" (page 125).

To delete a network interface, use the following command:

```
ibrix_nic -d -n IFNAME -h HOSTLIST
```

The following command deletes interface `eth3` from file serving nodes `s1.hp.com` and `s2.hp.com`:

```
ibrix_nic -d -n eth3 -h s1.hp.com,s2.hp.com
```

Viewing network interface information

Executing the `ibrix_nic` command with no arguments lists all interfaces on all file serving nodes. Include the `-h` option to list interfaces on specific hosts.

```
ibrix_nic -l -h HOSTLIST
```

The following table describes the fields in the output.

Field	Description
BACKUP HOST	File serving node for the standby network interface.
BACKUP-IF	Standby network interface.
HOST	File serving node.
IFNAME	Network interface on this file serving node.

Field	Description
IP_ADDRESS	IP address of this NIC.
LINKMON	Whether monitoring is on for this NIC.
MAC_ADDR	MAC address of this NIC.
ROUTE	IP address in routing table used by this NIC.
STATE	Network interface state.
TYPE	Network type (cluster or user).

When `ibrix_nic` is used with the `-i` option, it reports detailed information about the interfaces. Use the `-h` option to limit the output to specific hosts. Use the `-n` option to view information for a specific interface.

```
ibrix_nic -i [-h HOSTLIST] [-n NAME]
```

12 Licensing

This chapter describes how to view your current license terms and how to obtain and install new StoreAll software product license keys.

NOTE: For licensing features such as block snapshots on the HP P2000 G3 MSA Array System or HP 2000 Modular Smart Array, see the array documentation.

Viewing license terms

The StoreAll software license file is stored in the installation directory. To view the license from the GUI, select **Cluster Configuration** in the Navigator and then select **License**.

To view the license from the CLI, use the following command:

```
ibrix_license -i
```

The output reports your current node count and capacity limit. In the output, Segment Server refers to file serving nodes.

Retrieving a license key

When you purchased this product, you received a License Entitlement Certificate. You will need information from this certificate to retrieve and enter your license keys.

You can use any of the following methods to request a license key:

- Obtain a license key from http://webware.hp.com.
- Use AutoPass to retrieve and install permanent license keys. See "Using AutoPass to retrieve and install permanent license keys" (page 128).
- Fax the Password Request Form that came with your License Entitlement Certificate. See the certificate for fax numbers in your area.
- Call or email the HP Password Center. See the certificate for telephone numbers in your area or email addresses.

Using AutoPass to retrieve and install permanent license keys

The procedure must be run from a client with JRE 1.5 or later installed and with a desktop manager running (for example, a Linux-based system running X Windows). The `ssh` client must also be installed.

1. On the Linux-based system, run the following command to connect to the Fusion Manager:

   ```
   ssh -X root@<management_console_IP>
   ```

2. When prompted, enter the password for the Fusion Manager.
3. Launch the AutoPass GUI:

   ```
   /usr/local/ibrix/bin/fusion-license-manager
   ```

4. In the AutoPass GUI, go to **Tools**, select **Configure Proxy**, and configure your proxy settings.
5. Click **Retrieve/Install License** > **Key** and then retrieve and install your license key.

   If the Fusion Manager machine does not have an Internet connection, retrieve the license from a machine that does have a connection, deliver the file with the license to the Fusion Manager machine, and then use the AutoPass GUI to import the license.

13 Upgrading firmware

The Firmware Management Tool (FMT) is a utility that scans the StoreAll system for outdated firmware and provides a comprehensive report that provides the following information:

- Device found
- Active firmware found on the discovered device
- Qualified firmware for the discovered device
- Proposed action — Users are told whether an upgrade is recommended
- Severity — How severe an upgrade is required
- Reboot required on flash
- Device information
- Parent device ID

Components for firmware upgrades

The HP StoreAll system includes several components with upgradable firmware. The following table lists these components and specifies whether they can be upgraded online and in a nondisruptive manner. The following example is an example for the 9320 system.

Component	Online and Nondisruptive?
DL380	Nondisruptive if done one server at a time
SAS HBA	Yes, if done one server at a time
OS image	Yes, if done one server at a time
RAID controller	Yes, if done one controller at a time
HDD	No

Enter the following command to show which components could be flagged for flash upgrade.

```
hpsp_fmt -lc
```

The following is an example of the server components that are displayed:

```
SERVER:
------
ILO3              :Integrated Lights-Out(iLO) 3
BIOS              :Systems ROM for Server Blade
Power_Mgmt_Ctlr   :Power Management Controller
Smart_Array_Ctlr  :HP Embedded Smart Array Controller
NIC               :HP Embedded Network Adapter
PCIeNIC           :HP PCIe Network Adapter
SERVER_HDD        :HP Server Hard Disk Drives
```

Steps for upgrading the firmware

IMPORTANT: On the StoreAll 9320 Storage platform the prerequisite measures must be followed before performing the firmware update operation on enclosure hard disk drives:

- Storage disk drive update is an OFFLINE process. Ensure that all host and array I/O must be stopped prior to the update.
- Make sure all the file systems are unmounted. Failure to comply may result in an OS crashing.
- Ensure that no other user is performing administrative functions on the storage system.
- Ensure that the storage system configuration is not being reconfigured during a firmware update.
- To avoid spurious failure indications and potential system crashes, you must suspend hardware monitoring during the update process. You should execute `hpspmonControl --pause`on both servers in the couplet before performing the flash operation. You should then make sure to execute `hpspmonControl --resume` on both servers in the couplet after the flash operation is complete.
- Do not cycle power or restart devices during a firmware update.

To upgrade the firmware for components:

1. Run the `/opt/hp/platform/bin/hpsp_fmt -fr` command to verify that the firmware on this node and subsequent nodes in this cluster is correct and up-to-date. This command should be performed before placing the cluster back into service.

   The following figure shows an example of the firmware recommendation output and corrective component upgrade flash:

```
root@ib149-124:~
[root@ib149-124 ~]# hpsp_fmt -fr
Since no category argument is specified, operation will be performed on server
components...
Based on the configuration, Flash recommendation may take 5 minutes.
     <<< Completed >>>

-------------------------------------------------------------------------------------------------------------------------------------------------------
Device            Device ID           Active FW     Qualified     Proposed   Severity   Reboot Reqd.  Device Information   Parent Device ID
                                      Version       FW Version    Action                on Flash
-------------------------------------------------------------------------------------------------------------------------------------------------------
BIOS              USE049NCAB          P67           P67           UPGRADE    MANDATORY  Yes           N/A                 N/A
                                      (12/01/2010)  (05/05/2011)
iLO3              ILOUSE049NCAB       1.16          1.28          UPGRADE    MANDATORY  No            N/A                 N/A
NIC               eth0                BC:5.2.3      BC:5.2.3      UPGRADE    MANDATORY  Yes           NC382i              N/A
                                      (iSCSI:4.2.10) (iSCSI:7.2.53)
NIC               eth2                BC:5.2.3      BC:5.2.3      UPGRADE    MANDATORY  Yes           NC382i              N/A
                                      (iSCSI:4.2.10) (iSCSI:7.2.53)
PCIeNIC           N/A                 5.12-2        5.12-2        NONE       NONE       No            NC364T              N/A
Power_Mgmt_Ctlr   N/A                 1.6           1.6           NONE       NONE       Yes           N/A                 N/A
SERVER_HDD        PLVXUTTE            HPD5          HPD5          NONE       NONE       Yes           PORT:1I BOX:1 BAY:1 Ctlr:5001438010229BF0
SERVER_HDD        PLVW5O9E            HPD5          HPD5          NONE       NONE       Yes           PORT:1I BOX:1 BAY:2 Ctlr:5001438010229BF0
Smart_Array_Ctlr  5001438010229BF0    5.14          5.70          UPGRADE    MANDATORY  Yes           P410i               Chas:USE049NCAB
-------------------------------------------------------------------------------------------------------------------------------------------------------
```

IMPORTANT: For some components in StoreAll 9320/9300, performing a firmware update would require some prerequisite measures that must be followed before the flash operation. The Firmware update command, `hpsm_fmt -flashrec -c <component-name>`, provides the prerequisites required and it asks users they have taken necessary steps for the firmware update.

IMPORTANT: Upgrade the firmware in the following order:

1. Server
2. Chassis
3. Storage

2. Do the following based on the Proposed Action and Severity:

Status in Proposed Action column	Status in Severity column	Go to
UPGRADE	MANDATORY	Step 3
UPGRADE	RECOMMENDED	Step 3 is optional. However, it is recommended to perform step 3 for system stability and to avoid any known issues.
NONE or DOWNGRADE	MANDATORY	Step 4
NONE or DOWNGRADE	RECOMMENDED	Step 4 is optional. However, it is recommended to perform step 4 for system stability and to avoid any known issues.

3. Perform the flash operation by entering the following command and then go to step 5:

   `hpsp_fmt -flash -c <components-name>`

   The following screen shot displays a successful flash operation.

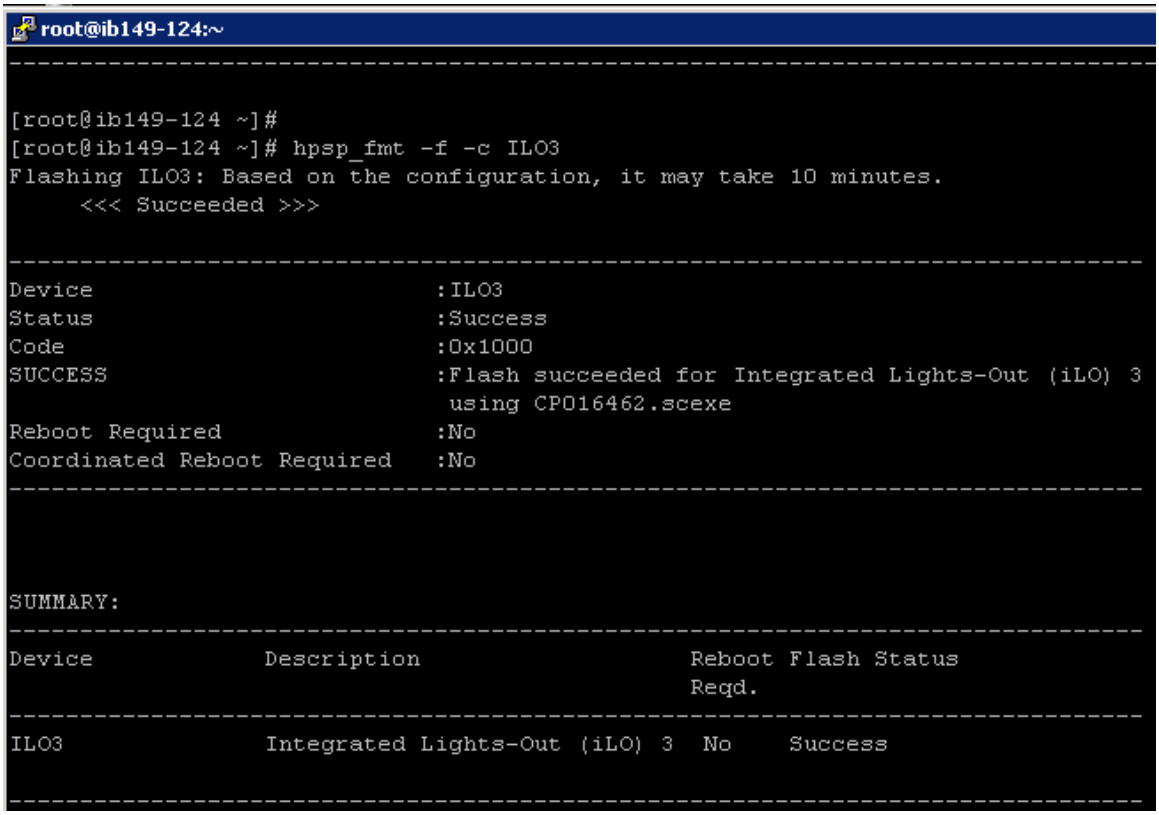
root@ib149-124:~
[root@ib149-124 ~]#
[root@ib149-124 ~]# hpsp_fmt -f -c ILO3
Flashing ILO3: Based on the configuration, it may take 10 minutes.
<<< Succeeded >>>
Device :ILO3
Status :Success
Code :0x1000
SUCCESS :Flash succeeded for Integrated Lights-Out (iLO) 3 using CP016462.scexe
Reboot Required :No
Coordinated Reboot Required :No
SUMMARY:
Device Description Reboot Reqd. Flash Status
ILO3 Integrated Lights-Out (iLO) 3 No Success

4. Perform the flash operation by entering the following command and then go to step 5:

   `hpsp_fmt -flash -c <components-name> --force`

5. If the components require a reboot on flash, failover the FSN for continuous operation as described in the following steps:

   NOTE: Although the following steps are based on a two-node cluster, all steps can be used in a multiple node clusters.

a. Determine whether the node to be flashed is the active Fusion Manager by enter the following command:

   ```
   ibrix_fm -i
   ```

b. Perform a manual FM failover on the local node by entering the following command from the active Fusion Manager:

   ```
   ibrix_fm -m nofmfailover server1
   ```

   The FM failover will take approximately one minute.

c. If server1 is not the active Fusion Manager, proceed to step e to fail over server1 to server2.

d. To see which node is now the active Fusion Manager, enter the following command:

   ```
   ibrix_fm -i
   ```

e. Move to your new active Fusion Manager node, and then enter the following command to perform the failover:

   ```
   ibrix_server -f -p -h server1
   ```

   NOTE: The `-p` switch in the failover operation lets you reboot the effected node and in turn the flash of the following components:

   - BIOS
   - NIC
   - PcIeNIC
   - Power_Mgmt_Ctlr
   - SERVER_HDD
   - Smart_Array_Ctlr

f. Once the FSN boots up, verify the software reports the FSN as `Up, FailedOver` by enter the following command:

   ```
   ibrix_server -l
   ```

g. Confirm the recommended flash was completed successfully by enter the following command:

   ```
   hpsp_fmt -fr server -o /tmp/fwrecommend.out
   ```

   Verify that the `Proposed Action` column requires no more actions, and the `Active FW Version` and `Qualified FW Version` columns display the same values.

h. Fail back your updated server by entering the following command:

   ```
   ibrix_server -f -U -h server1
   ```

i. The failed-over Fusion Manager remains in nofmfailover mode until it is moved to passive mode by using the following command:

   ```
   ibrix_fm -m passive
   ```

   NOTE: A Fusion Manager cannot be moved from nofmfailover mode to active mode.

j. Repeat steps a through h for the backup server component, in this example switch server1 with server2 in the commands.

k. Repeat steps a through I for each node that requires a firmware upgrade.

6. If you are upgrading to 6.3, you must complete the steps provided in the "After the upgrade" section for your type of upgrade, as shown in the following table:

Type of upgrade	Complete the steps in this section
Online upgrades	"After the upgrade" (page 13)
Automated offline upgrades	"After the upgrade" (page 15)
Manual offline upgrades	"After the upgrade" (page 17)

Finding additional information on FMT

You can find additional information on FMT as follows:

- Online help for FMT. To access the online help for FMT, enter the `hpsp_fmt` command on the file system node console.
- *HP HPSP_FMT User Guide*. To access the *HP HPSP_FMT User Guide*, go to the HP StoreAll Storage Manuals page:

  http://www.hp.com/support/StoreAllManuals

Downloading MSA2000 G2/G3 firmware for 9320 systems

To obtain the firmware, complete the following steps:

1. Go to the following HP web site:

   http://www.hp.com/go/support
2. Select **Download drivers and software** , enter your model number in the search box, and begin your search.
3. Select your array type.
4. Select your operating system.
5. Find the firmware version listed in the table, and select the hyperlink in the far left column.
6. To download the installation instructions, select the **Installation Instructions** tab.
7. To download the firmware, click **Download**.

14 Troubleshooting

Collecting information for HP Support with the IbrixCollect

Ibrix Collect is a log collection utility that allows you collect relevant information for diagnosis by HP Support when system issues occur. The collection can be triggered manually using the GUI or CLI, or automatically during a system crash. Ibrix Collect gathers the following information:

- Specific operating system and StoreAll command results and logs
- Crash digester results
- Summary of collected logs including error/exception/failure messages
- Collection of information from LHN and MSA storage connected to the cluster

NOTE: When the cluster is upgraded from a StoreAll software version earlier than 6.0, the support tickets collected using the `ibrix_supportticket` command will be deleted. Before performing the upgrade, download a copy of the archive files (`.tgz`) from the `/admin/platform/diag/supporttickets` directory.

Collecting logs

To collect logs and command results using the GUI:

1. Select **Cluster Configuration**, and then select **Data Collection**.
2. Click **Collect**.

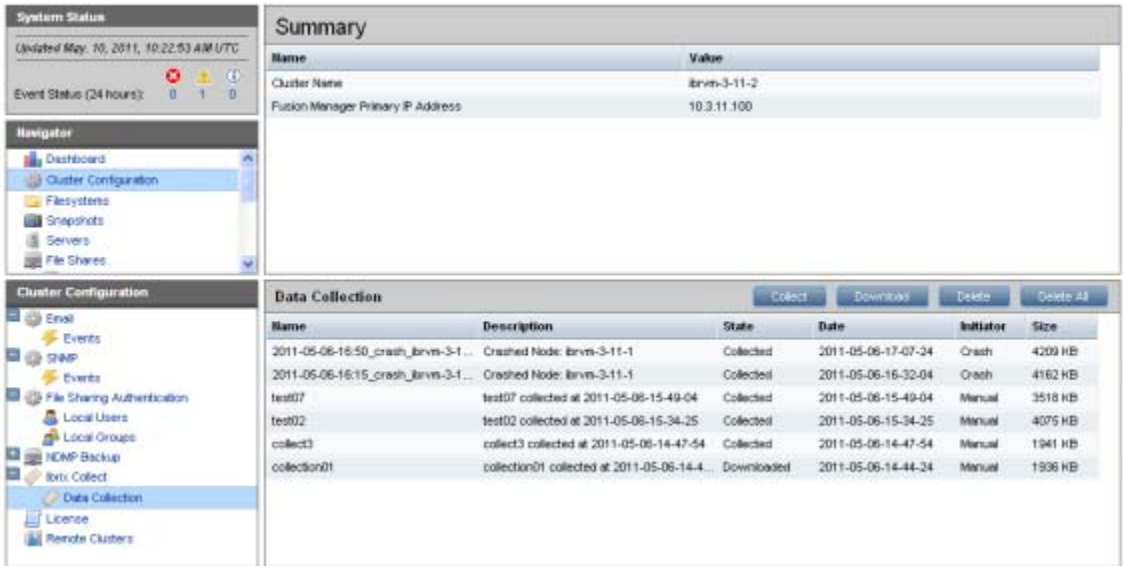

3. The data is stored locally on each node in a compressed archive file `<nodename>_<filename>_<timestamp>.tgz` under `/local/ibrixcollect`.

   Enter the name of the `.tgz` file that contains the collected data. The default location to store this `.tgz` file is located on the active Fusion Manager node at `/local/ibrixcollect/archive`.

Collect Data
* Filename:
Save in: /local/ibrixcollect/archive
Please provide the name for the zip file which will contain the collected logs and command outputs. The zip file will be saved in the default location, /local/ibrixcollect/archive in the active fusion manager.
(*) Required Value
OK
Cancel
Help

4. Click **Okay**.

To collect logs and command results using the CLI, use the following command:

```
ibrix_collect -c -n NAME
```

NOTE: Only one manual collection of data is allowed at a time.

NOTE: When a node restores from a system crash, the `vmcore` under `/var/crash/<timestamp>` directory is processed. Once processed, the directory will be renamed `/var/crash/<timestamp>_PROCESSED`. HP Support may request that you send this information to assist in resolving the system crash.

NOTE: HP recommends that you maintain your crash dumps in the `/var/crash` directory. Ibrix Collect processes the core dumps present in the `/var/crash` directory (linked to `/local/platform/crash`) only. HP also recommends that you monitor this directory and remove unnecessary processed crashes.

Downloading the archive file

When data is collected, a compressed archive file is created and stored in a compressed archive file (`.tgz`) under `/local/ibrixcollect/archive` directory. To download the collected data to your desktop, select the collection and click **Download** from the Fusion Manager.

NOTE: Only one collection can be downloaded at a time.

NOTE: The average size of the archive file depends on the size of the logs present on individual nodes in the cluster.

NOTE: You may later be asked to email this final `.tgz` file to HP Support.

Deleting the archive file

You can delete a specific data collection or all collections simultaneously in the GUI and the CLI.

To delete a specific data collection using the GUI, select the collection to be deleted, and click **Delete**. The `tgz` file stored locally will be deleted from each node.

To delete all of the collections, click **Delete All**.

To delete a specific data collection using the CLI, use the following command:

```
ibrix_collect -d -n NAME
```

To specify more than one collection to be deleted at a time from the CLI, provide the names separated by a semicolon.

To delete all data collections manually from the CLI, use the following command:

```
ibrix_collect -d -F
```

Configuring Ibrix Collect

You can configure data collection to occur automatically upon a system crash. This collection will include additional crash digester output. The archive filename of the system crash-triggered collection will be in the format `<timestamp>_crash_<crashedNodeName>.tgz`.

1. To enable or disable an automatic collection of data after a system crash, and to configure the number of data sets to be retained:
   a. Select **Cluster Configuration**, and then select **Ibrix Collect**.
   b. Click **Modify**, and the following dialog box will appear.

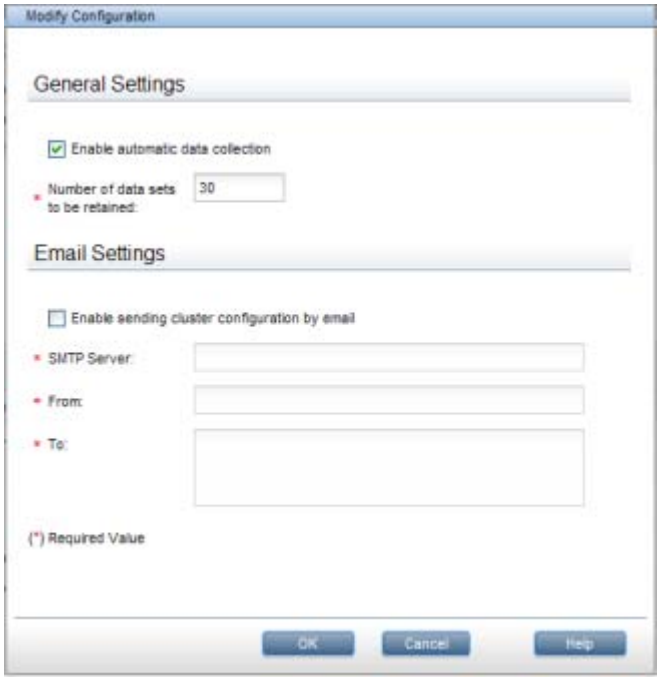

   c. Under General Settings, enable or disable automatic collection by checking or unchecking the appropriate box.
   d. Enter the number of data sets to be retained in the cluster in the text box.

   To enable/disable automatic data collection using the CLI, use the following command:

   ```
   ibrix_collect -C -a <Yes\No>
   ```

   To set the number of data sets to be retained in the cluster using the CLI, use the following command:

   ```
   ibrix_collect -C -r NUMBER
   ```

2. To configure emails containing a summary of collected information of each node to be sent automatically to your desktop after every data collection event:
   a. Select **Cluster Configuration**, and then select **Ibrix Collect**.
   b. Click **Modify**.
   c. Under Email Settings, enable or disable sending cluster configuration by email by checking or unchecking the appropriate box.
   d. Fill in the remaining required fields for the cluster configuration and click **Okay**.

To set up email settings to send cluster configurations using the CLI, use the following command:

```
ibrix_collect -C -m <Yes\No> [-s <SMTP_server>] [-f <From>] [-t <To>]
```

NOTE: More than one email ID can be specified for `-t` option, separated by a semicolon. The "From" and "To" command for this SMTP server are Ibrix Collect specific.

Obtaining custom logging from `ibrix_collect` add-on scripts

You can create add-on scripts that capture custom StoreAll and operating system commands and logs. To activate an add-on script, place it in the specified location, and the add-on script will run when the `ibrix_collect` command is executed. Output of these add-on scripts is packaged into IbrixCollect tar files.

Table 3 `ibrix_collect` add-on scripts

Step	Description	Where to find more information?
1	Create an add-on script.	"Creating an add-on script" (page 137)
2	Run the add-on script.	"Running an add-on script" (page 138)
3	View the output from the add-on script.	"Viewing the output from an add-on script" (page 138)

Creating an add-on script

To create an add-on script:

1. Add-on Scripts names should be in the defined format. The file name of the script should strictly followed this format `<release Number>_<add-on_script_name>.sh`.

   When you provide the release number in the file name, remove the period between the first and second digit of the version number.

   For example, assume you obtained the version number from the `ibrix_version -l` command.

   The output of the command displays the version number.

   ```
   [root@host2 /]# ibrix_version -l
   Fusion Manager version: 6.3.33
   ==============================
   Segment Servers
   ===============
   HOST_NAME FILE_SYSTEM IAD/IAS IAD/FS OS KERNEL_VERSION ARCH
   --------—— ------------------------------- ------- ------ ---------
   -------------- ----
   host2 6.3.33(internal rev 132818 in SVN) 6.3.33 6.3.33 GNU/Linux
   2.6.18-194.el5 x86_64
   ```

   In this example, the version displayed is 6.3.33. Use the first two digit of the version number, 6.3 for example, as a prefix to "Add-on Script File Name" without "." (Dot), so that the name of an add-on script named `AddOnTest.sh` would be `63_AddOnTest.sh`.

IMPORTANT: The version provided in the name must match the version of StoreAll on which you plan to run the script. For example, any add-on script that you want to run on StoreAll version 6.3, must have `63_` in its file name; otherwise, the script will not run. For example, if you prefix the name with another version, such as `6.2_`, and you attempt to run the add-on script on StoreAll 6.3, the script will not run.

2. Place the added-on script in the following directory:

   `/usr/local/ibrix/ibrixcollect/ibrix_collect_add_on_scripts/`

The following example shows several add-on scripts stored in the `ibrix_collect_add_on_scripts` directory:

```
root@host2 /]# ls -l
/usr/local/ibrix/ibrixcollect/ibrix_collect_add_on_scripts/
total 8
-rwxr-xr-x 1 root root 93 Dec 7 13:39 60_addOn.sh
-rwxrwxrwx 1 root root 48 Dec 20 09:22 63_AddOnTest.sh
```

3. Write an add-on shell script that contains a custom command/log that needs to be collected in the final StoreAll collection. Only StoreAll and operating system commands are supported in the scripts. These scripts should have appropriate permission to be executed.

IMPORTANT: Make sure the scripts that you are creating do not collect information or logs that are already collected as part of the `ibrix_collect` command.

4. Make sure that add-on scripts that collect the custom logs redirects the collected custom logs to the directory `/local/ibrixcollect/ibrix_collect_additional_data`. Only files copied to this location will be included in the generated IbrixCollect tar file. Output of the add-on scripts is available only when the IbrixCollect process is completed and the tar files containing the output are extracted. See "Running an add-on script" (page 138) and then "Viewing the output from an add-on script" (page 138).

Running an add-on script

To run an add-on script:

1. Verify that the add-on script is saved under the following location:

   `/usr/local/ibrix/ibrixcollect/ibrix_collect_add_on_scripts/`

   The `ibrix_collect` command only runs add-on scripts saved in this location.

2. Enter the `ibrix_collect` command:

   `ibrix_collect -c -n addOnCollection`

   In this instance *`addOnCollection`* is the collection name.

   The output of the add-on scripts is included into the final tar collection along with other logs and command outputs. In this instance, the output would be in the `addOnCollection.tgz` file.

NOTE: The add-on scripts timeout after 20 minutes.

Viewing the output from an add-on script

To view an output from an add-on script:

1. Go to the active Fusion Manager node in the `/local/ibrixcollect/archive` directory by entering the following command:

   `[root@host2 /]#cd /local/ibrixcollect/archive/`

2. The output of the add-on scripts is available under the tar file of the individual node. To view the contents of the directory, enter the following command:

   `[root@host2 /]#ls -l`

   The following is an example of the output displayed:

```
total 3520
-rw-r--r-- 1 root root 2021895 Dec 20 12:41 addOnCollection.tgz
```

3. Extract the tar file, containing the output of the add-on script. The tar file containing the output of the add-on script has the name of the collection by entering the following command:

   ```
   [root@host2 archive]#tar -xvf addOnCollection.tgz
   ```

   In this instance, `addOnCollection.tgz` is the tar file containing the output of the add-on script.

   The tar command displays the following:

   ```
   ./
   ./host2_addOnCollection_2012-12-20-12-38-36.tgz
   ```

4. Individual node files in the tar format are provided as *<hostname>_<collection-name>_<time-date stamp>*.tgz

   Extract the <hostname>_<collection-name>_<time-date stamp>.tgz tar file by entering the following command:

   ```
   [root@host2 archive]#tar -xvf
   host2_addOnCollection_2012-12-20-12-38-36.tgz
   ```

   In this instance, *host2_addOnCollection_2012-12-20-12-38-36.tgz* is the individual node file (<hostname>_<collection-name>_<time-date stamp>.tgz).

5. A directory with the host name is extracted. The output of the add-on script is found in the /<hostname>/logs/add_on_script/local/ibrixcollect/ibrix_collect_additional_data

   Find the directory containing the host name by entering the `ls -l` command, as shown in the following example:

   ```
   [root@host2 archive]# ls -l
   ```

   The following is the output of the command:

   ```
   total 5636
   -rw-r--r-- 1 root root 2021895 Dec 20 12:41 addOnCollection.tgz
   drwxr-xr-x 6 root root 4096 Dec 20 12:41 host2
   -rw-r--r-- 1 root root 2156388 Dec 20 12:41
   host2_addOnCollection_2012-12-20-12-38-36.tgz
   ```

   In this example, *host2* is the directory with the host name.

6. Go to the /<hostname>/logs/add_on_script/local/ibrixcollect/ibrix_collect_additional_data directory, which contains the output of the add-on script.

   ```
   [root@host2 archive]#cd
   host2/logs/add_on_script/local/ibrixcollect/ibrix_collect_additional_data/
   ```

   In this instance, *host2* is the name of the host.

7. View the contents of the /<hostname>/logs/add_on_script/local/ibrixcollect/ibrix_collect_additional_data directory:

   ```
   [root@host2 ibrix_collect_additional_data]#ls -l
   ```

   The command displays the following output:

   ```
   total 4
   -rw-r--r-- 1 root root 2636 Dec 20 12:39 63_AddOnTest.out
   ```

   In this instance, *63_AddOnTest.out* displays the output of the add-on script.

Viewing data collection information

To view data collection history from the CLI, use the following command:

```
ibrix_collect -l
```

To view data collection details such as date (of creation), size, description, state and initiator, use the following command:

```
ibrix_collect -v -n <Name>
```

Adding/deleting commands or logs in the XML file

To add or change the logs that are collected or commands that are executed during data collection, you can modify the Ibrix Collect XML files that are stored in the directory `/usr/local/ibrix/ibrixcollect`.

The `/usr/local/ibrix/ibrixcollect` commands executed and the logs collected during data collection are maintained in the following files under `/usr/local/ibrix/ibrixcollect` directory:

- `fm_summary.xml` – Commands pertaining to the Fusion Manager node
- `ss_summary.xml` – Commands pertaining to the file serving node
- `common_summary.xml` – Commands and logs common to both Fusion Manager and file serving nodes

NOTE: These xml files should be modified carefully. Any missing tags during modification might cause Ibrix Collect to not work properly.

Viewing software version numbers

To view version information for a list of hosts, use the following command:

```
ibrix_version -l [-h HOSTLIST]
```

For each host, the output includes:

- Version number of the installed file system
- Version numbers of the IAD and File System module
- Operating system type and OS kernel version
- Processor architecture

The `-S` option shows this information for all file serving nodes. The `-C` option shows the information for all StoreAll clients.

The file system and IAD/FS output fields should show matching version numbers unless you have installed special releases or patches. If the output fields show mismatched version numbers and you do not know of any reason for the mismatch, contact HP Support. A mismatch might affect the operation of your cluster.

Troubleshooting specific issues

Software services

Cannot start services on a file serving node, or Linux StoreAll client

SELinux might be enabled. To determine the current state of SELinux, use the `getenforce` command. If it returns `enforcing`, disable SELinux using either of these commands:

```
setenforce Permissive
setenforce 0
```

To permanently disable SELinux, edit its configuration file (`/etc/selinux/config`) and set `SELINUX=parameter` to either `permissive` or `disabled`. SELinux will be stopped at the next boot.

For StoreAll clients, the client might not be registered with the Fusion Manager. For information on registering clients, see the *HP StoreAll Storage Installation Guide*.

Failover

Cannot fail back from failover caused by storage subsystem failure

When a storage subsystem fails and automated failover is turned on, the Fusion Manager will initiate its failover protocol. It updates the configuration database to record that segment ownership has transferred from primary servers to their standbys and then attempts to migrate the segments to the standbys. However, segments cannot migrate because neither the primary servers nor the standbys can access the storage subsystem and the failover is stopped.

Perform the following manual recovery procedure:

1. Restore the failed storage subsystem (for example, replace failed Fibre Channel switches or replace a LUN that was removed from the storage array).
2. Reboot the standby servers, which will allow the failover to complete.

Cannot fail back because of a storage subsystem failure

This issue is similar to the previous issue. If a storage subsystem fails after you have initiated a failback, the configuration database will record that the failback occurred, even though segments never migrated back to the primary server. If you execute `ibrix_fs -i -f FSNAME`, the output will list `No` in the ONBACKUP field, indicating that the primary server now owns the segments, even though it does not. In this situation, you will be unable to complete the failback after you fix the storage subsystem problem.

Perform the following manual recovery procedure:

1. Restore the failed storage subsystem.
2. Reboot the primary server, which will allow the arrested failback to complete.

StoreAll client I/O errors following segment migration

Following successful segment migration to a different file serving node, the Fusion Manager sends all StoreAll clients an updated map reflecting the changes, which enables the clients to continue I/O operations. If, however, the network connection between a client and the Fusion Manager is not active, the client cannot receive the updated map, resulting in client I/O errors.

To fix the problem, restore the network connection between the clients and the Fusion Manager.

Windows StoreAll clients

Logged in but getting a “Permission Denied” message

The StoreAll client cannot access the Active Directory server because the domain name was not specified. Reconfigure the Active Directory settings, specifying the domain name. See the *HP StoreAll Storage Installation Guide* for more information.

Verify button in the Active Directory Settings tab does not work

This issue has the same cause as the above issue.

Mounted drive does not appear in Windows Explorer

To make a drive appear in Explorer, after mounting it, log off and then log back on, or reboot the machine. You can also open a MS-DOS command window and access the drive manually.

Mounted drive not visible when using Terminal Server

Refresh the browser's view of the system by logging off and then logging back on.

StoreAll client auto-startup interferes with debugging

The StoreAll client is set to start automatically, which can interfere with debugging a Windows StoreAll client problem. To prevent this, reboot the machine in safe mode and change the Windows StoreAll client service mode to manual, which enables you to reboot without starting the client.

1. Open the Services control manager (**Control Panel** > **Administrative Tools** > **Services**).
2. Right-click **StoreAll client Services** and select **Properties**.
3. Change the startup type to **Manual**, and then click **OK**.
4. Debug the client problem. When finished, switch the Windows StoreAll client service back to automatic startup at boot time by repeating these steps and changing the startup type to **Automatic**.

Synchronizing information on file serving nodes and the configuration database

To maintain access to a file system, file serving nodes must have current information about the file system. HP recommends that you execute `ibrix_health` on a regular basis to monitor the health of this information. If the information becomes outdated on a file serving node, execute `ibrix_dbck -o` to resynchronize the server's information with the configuration database. For information on `ibrix_health`, see "Monitoring cluster health" (page 95).

NOTE: The `ibrix_dbck` command should be used only under the direction of HP Support.

To run a health check on a file serving node, use the following command:

```
ibrix_health -i -h HOSTLIST
```

If the last line of the output reports `Passed`, the file system information on the file serving node and Fusion Manager is consistent.

To repair file serving node information, use the following command:

```
ibrix_dbck -o -f FSNAME [-h HOSTLIST]
```

To repair information on all file serving nodes, omit the `-h HOSTLIST` argument.

Troubleshooting an Express Query Manual Intervention Failure (MIF)

An Express Query Manual Intervention Failure (MIF) is a critical error that occurred during Express Query execution. These are failures Express Query cannot recover from automatically. After a MIF occurrence, the specific file system is logically removed from the Express Query and it requires a manual intervention to perform the recovery. Although these errors inhibit the normal functionality of the Express Query, they are typically due to another unrelated event in the cluster or the file system. Therefore, most of the work to recover from an Express Query MIF is to check the health of the cluster and the file system and take corrective actions to fix the issues caused by these events. Once the cluster and file system have an OK status, the MIF status can be cleared since the Express Query service will be recovering and restarting automatically.

In some very rare cases, a database corruption might occur, as a result of these external events or from some internal dysfunction. Express Query contains a recovery mechanism that tries to rebuild the database from information Express Query is keeping specifically for that critical situation. Express Query might be unable to recover from internal database corruption. Even though it is unlikely, it is possible and it might occur in the following two cases:

- A corrupted database has to be rebuilt from data that has been already backed up. If the data needed has been backed up, there is no automated way for Express Query to recover since that information has been deleted from the StoreAll file system after the backup. It is however possible to replay the database logs from the backup.
- Some data needed to rebuild the database is corrupted and therefore it cannot be used.

Even though database files, as well as information used in database recovery are well protected against corruption, corruption occurrence might occur.

NOTE: When a file system is in the MIF state, Express Query event recording is still occurring. When the database is re-enabled, the recorded events are processed and the database is synchronized with the file system again.

To recover from an Express Query Manual Intervention Failure (MIF):

1. Check the health of the file system as described in the "Monitoring cluster operations" (page 84), and clear any pending issues related to the file system (`<FSNAME>`).
2. Clear the Express Query MIF state by entering the following command:

   ```
   ibrix_archiving -C <FSNAME>
   ```

3. Monitor the Express Query recovery by entering the following command:

   ```
   ibrix_archiving -l
   ```

   While the Express Query is recovering from MIF, it displays the RECOVERY state. Wait for the state to return to `OK` or `MIF`.

   If the state returns as `OK`, no additional steps are required. The Express Query is updating the database with all the outstanding logged file system changes since the MIF occurrence.
4. If you have a MIF condition for one or several file systems and cluster and file system health checks are not `OK`, redo the previous steps.
5. Cluster and file system health checks have an `OK` status but Express Query is yet in a `MIF` condition for one or several specific file systems. This unlikely situation occurs when some data has been corrupted and it cannot be recovered.

   To solve this situation:

   a. If there is a full backup of the file system involved, do a restore.
   b. If there is no full backup:
      1. Disable Express Query for the file system, by entering the following command:

         ```
         ibrix_fs -T -D -f <FSNAME>
         ```

      2. Delete the current database for the file system by entering the following command:

         ```
         rm -rf <FS_MOUNTPOINT>/.archiving/database
         ```

      3. Enable Express Query for the file system by entering the following command:

         ```
         ibrix_fs -T -E -f <FSNAME>
         ```

         NOTE: The moment Express Query is enabled, database repopulation starts for the file system specified by `<FSNAME>`.

      4. If there are any backup of the Custom Metadata with the tool MDExport, re-import them with MDImport as described in *CLI Reference Guide*.

         NOTE: If no such backup exists, the custom metadata must be manually created again.

      5. Wait for the resynchronizer to complete by entering the following command:

         ```
         ibrix_archiving -l
         ```

         Repeat this command until it displays the `OK` status for the file system.
      6. If none of the above worked, contact HP.

15 Recovering a file serving node

Use the following procedure to recover a failed file serving node. You will need to create a QuickRestore DVD or USB key, as described later, and then install it on the affected node. This step installs the operating system and StoreAll software on the node and launches a configuration wizard.

△ **CAUTION:** The Quick Restore DVD or USB key restores the file serving node to its original factory state. This is a destructive process that completely erases all of the data on local hard drives.

Obtaining the latest StoreAll software release

StoreAll OS version 6.3 is only available through the registered release process. To obtain the ISO image, contact HP Support to register for the release and obtain access to the software dropbox.

Use a DVD

1. Burn the ISO image to a DVD.
2. Insert the DVD in the server.
3. Restart the server to boot from the DVD.
4. When the HP Network Storage System screen appears, enter **qr** to install the software.

Use a USB key

1. Copy the ISO to a Linux system.
2. Insert a USB key into the Linux system.
3. Execute `cat /proc/partitions` to find the USB device partition, which is displayed as `dev/sdX`. For example:

   ```
   cat /proc/partitions
   major minor  #blocks   name
   8     128    15633408  sdi
   ```

4. Execute the following `dd` command to make USB the QR installer:

   `dd if=<ISO file name with path> of=/dev/sdi oflag=direct bs=1M`

   For example:

   ```
   dd if=X9000-QRDVD-6.3.72-1.x86_64.signed.iso of=/dev/sdi oflag=direct bs=1M
   4491+0 records in
   4491+0 records out
   4709154816 bytes (4.7 GB) copied, 957.784 seconds, 4.9 MB/s
   ```

5. Insert the USB key into the server.
6. Boot the server from USB key. (Press **F11** and use option **3**.)
7. When the Network Storage System screen appears, enter **qr** to install the software.

Performing the recovery

Complete these steps:

1. Log into the node.
2. On the Individual Server Setup dialog box, enter your node specific information and click **OK**.

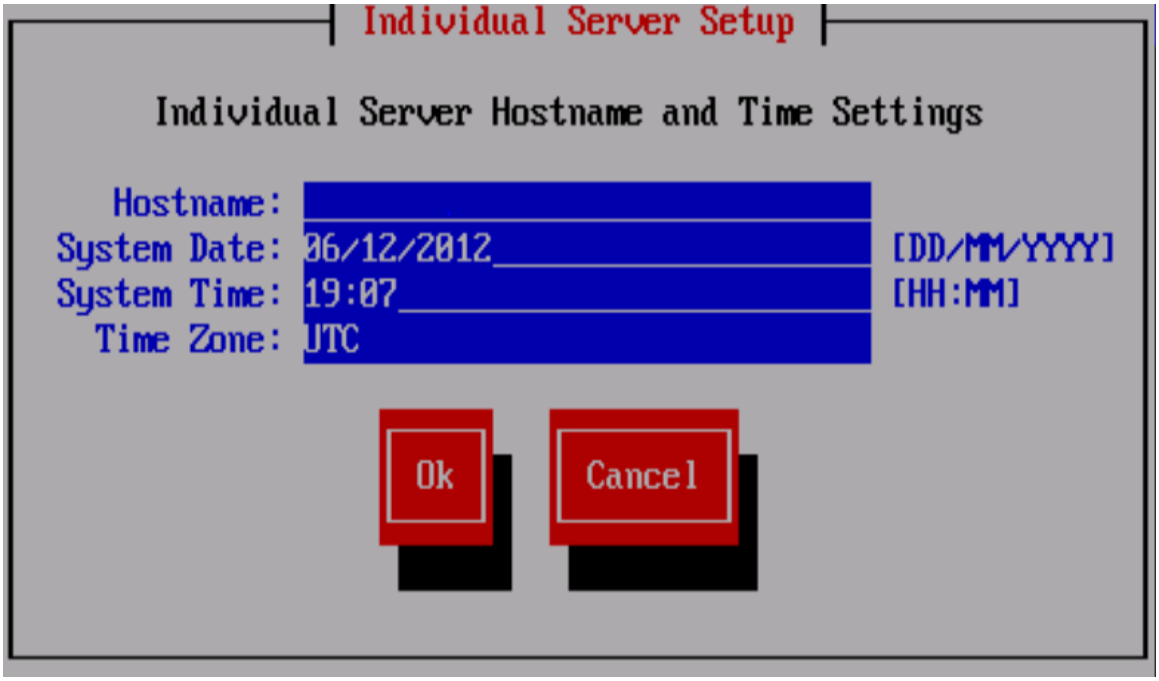

3. Enter the information for the node being restored on the Network Configuration dialog box and click **OK**.

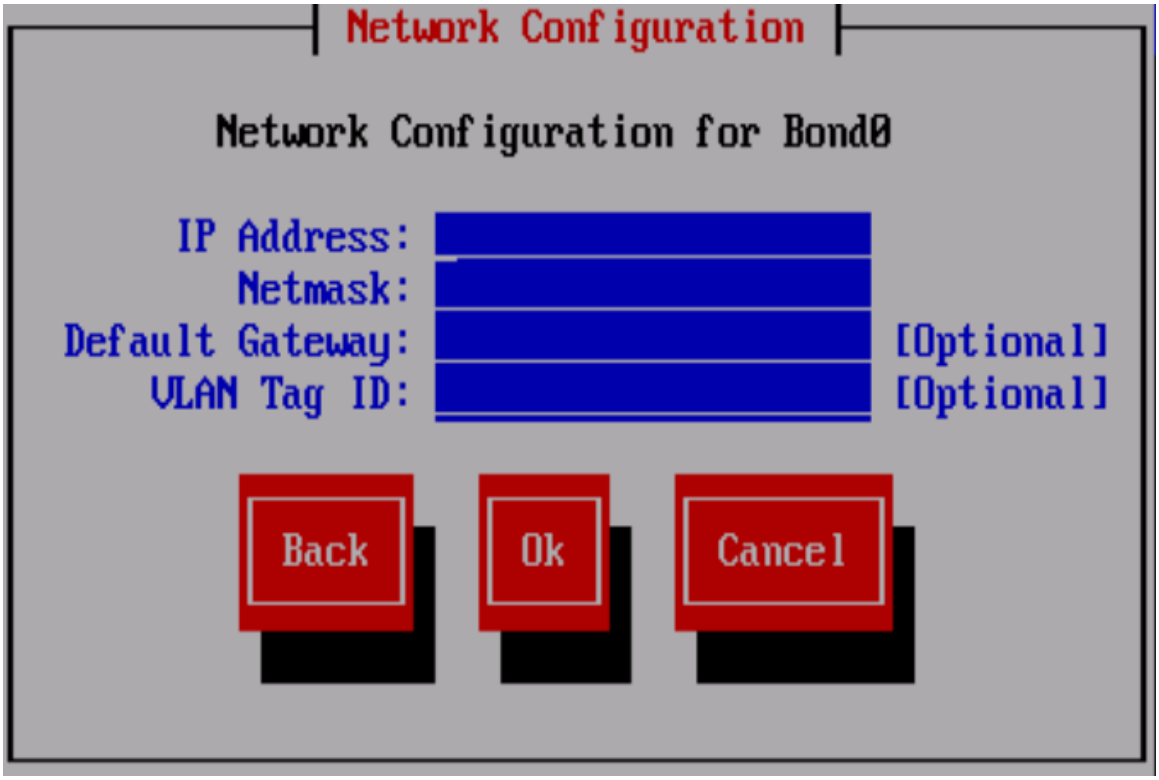

4. Confirm that the information displayed in the Configuration Summary dialog box is correct and click **Commit**.

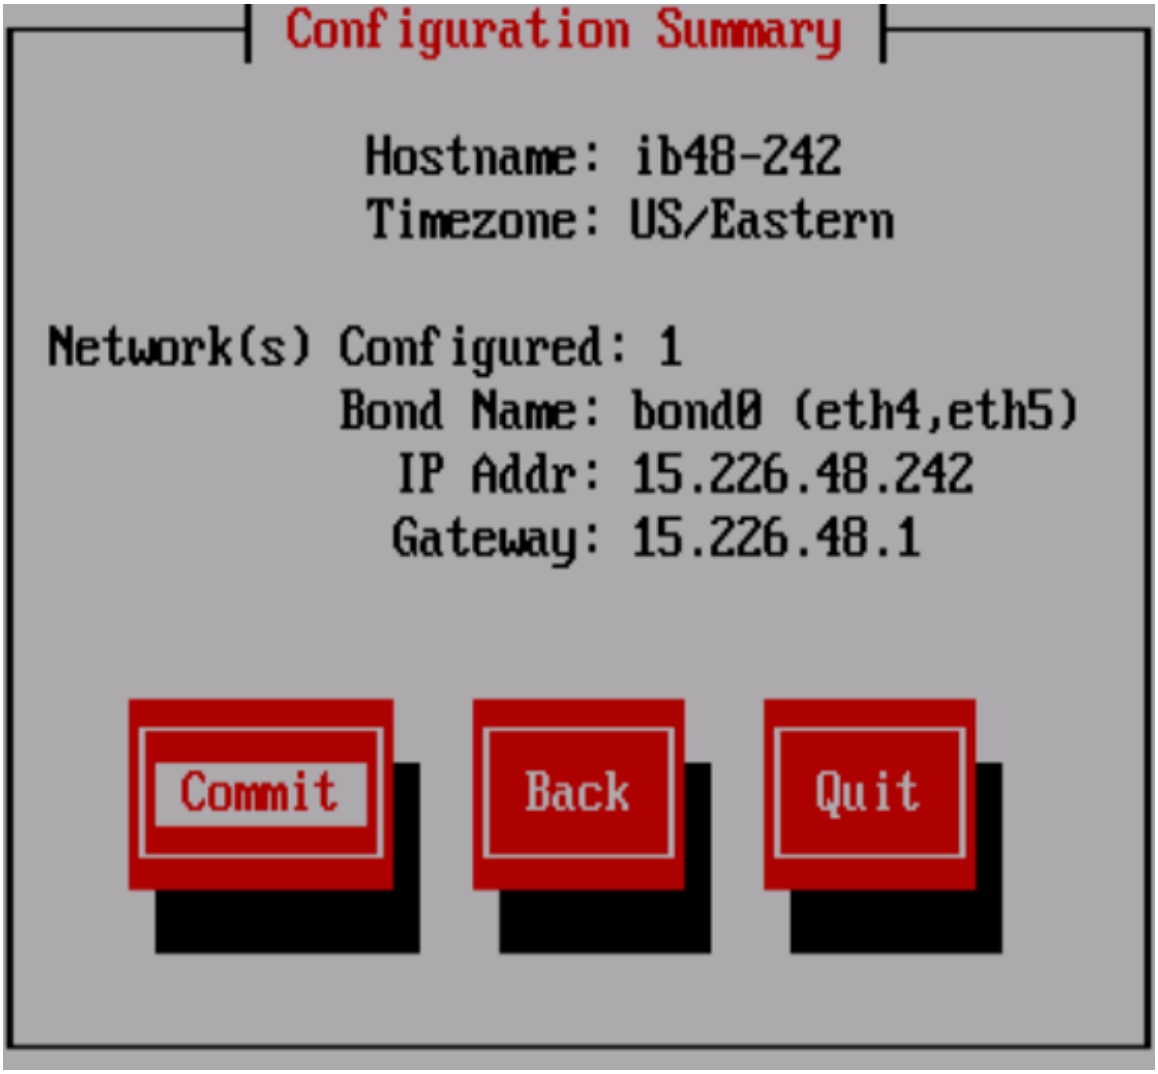

5. On the X9000 Installation — Network Setup Complete dialog box, select **Join this IBIRX server to an existing cluster** and click **OK**.

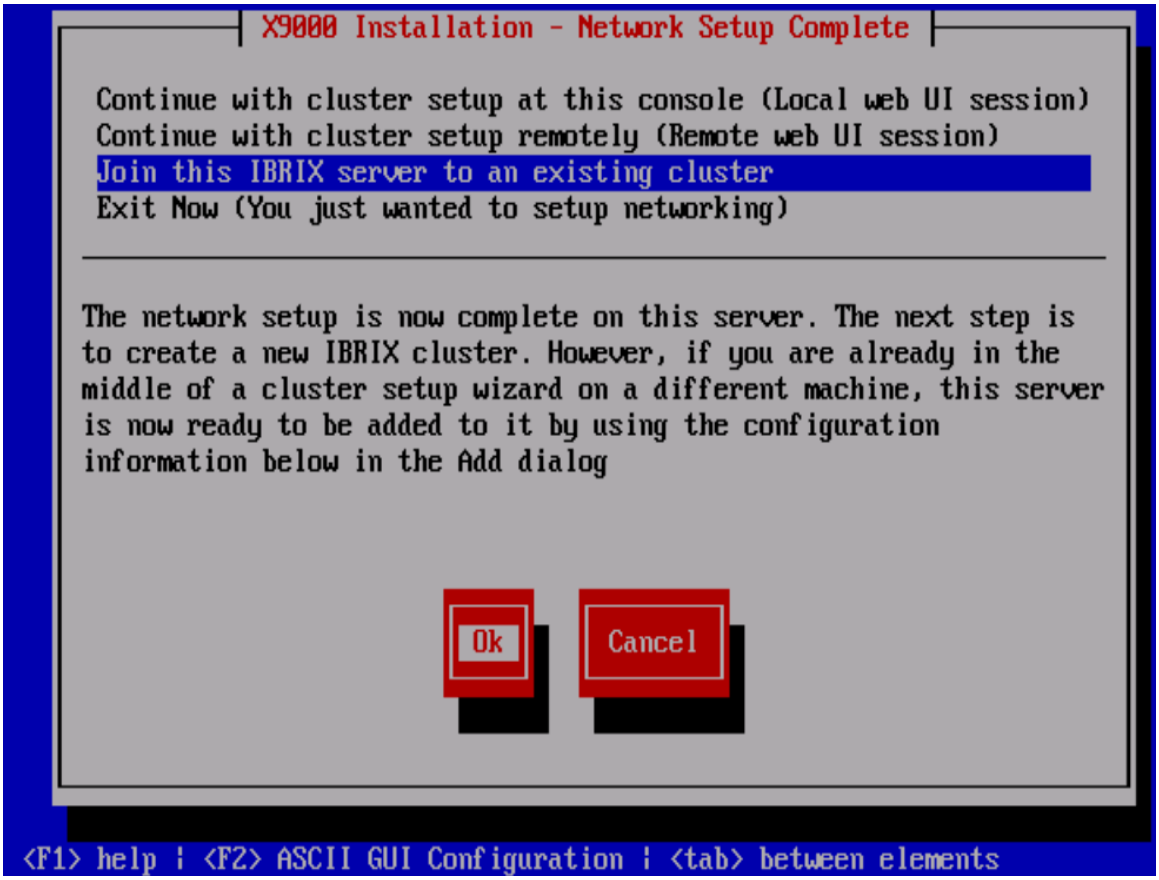
X9000 Installation - Network Setup Complete
Continue with cluster setup at this console (Local web UI session)
Continue with cluster setup remotely (Remote web UI session)
Join this IBRIX server to an existing cluster
Exit Now (You just wanted to setup networking)
The network setup is now complete on this server. The next step is to create a new IBRIX cluster. However, if you are already in the middle of a cluster setup wizard on a different machine, this server is now ready to be added to it by using the configuration information below in the Add dialog
Ok
Cancel
<F1> help ¦ <F2> ASCII GUI Configuration ¦ <tab> between elements

6. The wizard scans the network for existing clusters. On the Join Cluster dialog box, select the management console (Fusion Manager) for your cluster, and then click **OK**. If your cluster does not exist in the list of choices, click **Cancel** so that you can provide the IP address of the FM to which this node has to be registered.

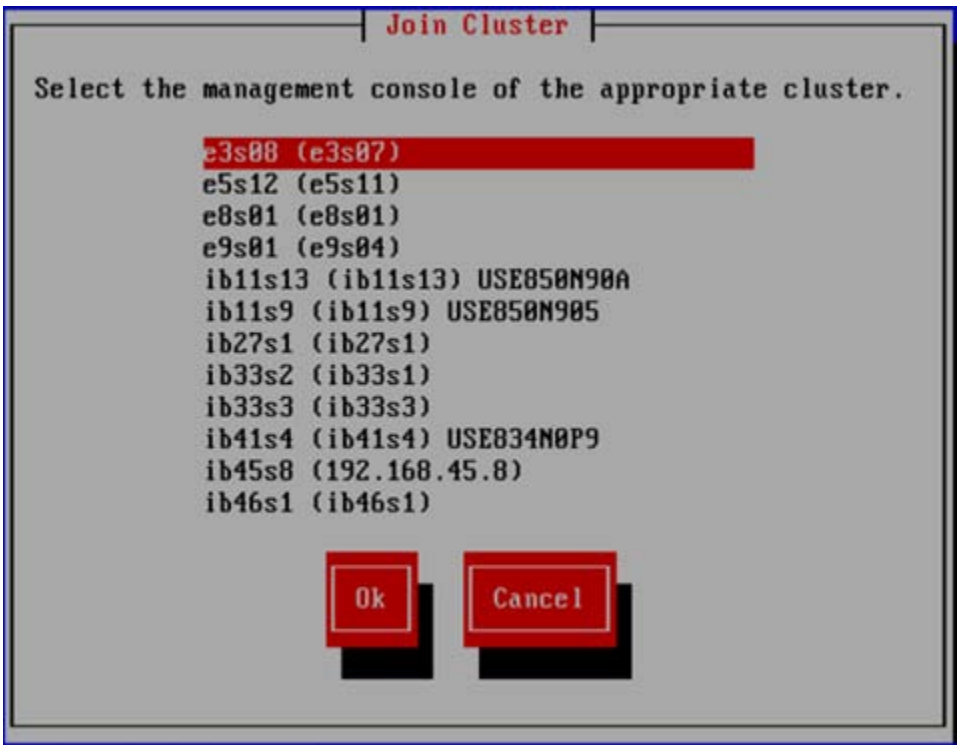

7. If you clicked the **Cancel** button in the previous dialog box, enter the management console IP of the desired cluster on the Management Console IP dialog box and click **OK**.

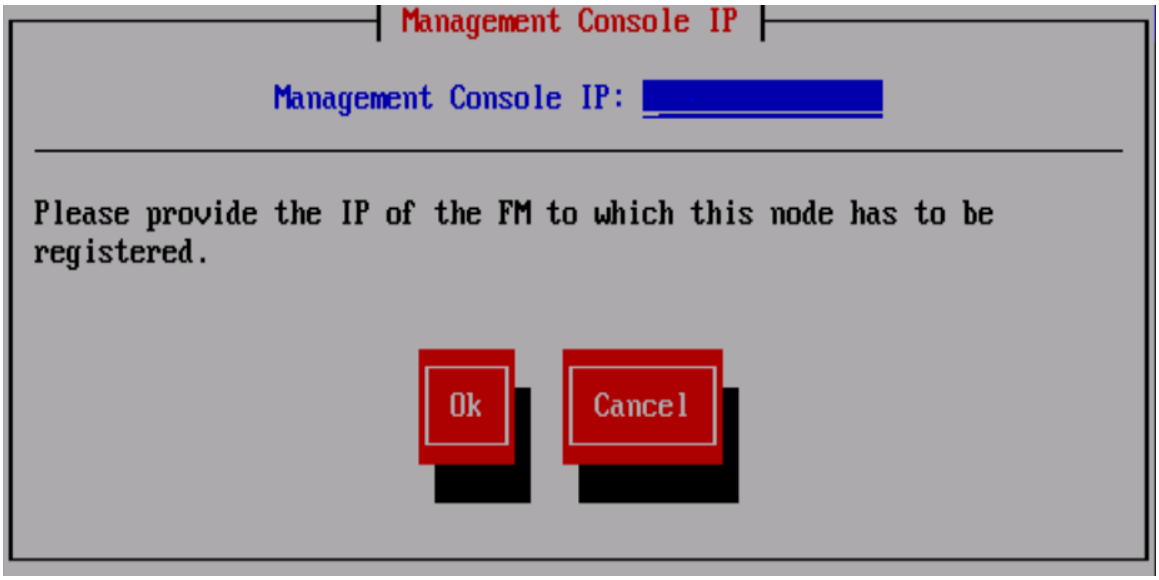
Management Console IP
Management Console IP:
Please provide the IP of the FM to which this node has to be registered.
Ok
Cancel

8. On the Replace Existing Server dialog box, click **Yes** when you are asked if you want to replace the existing server.

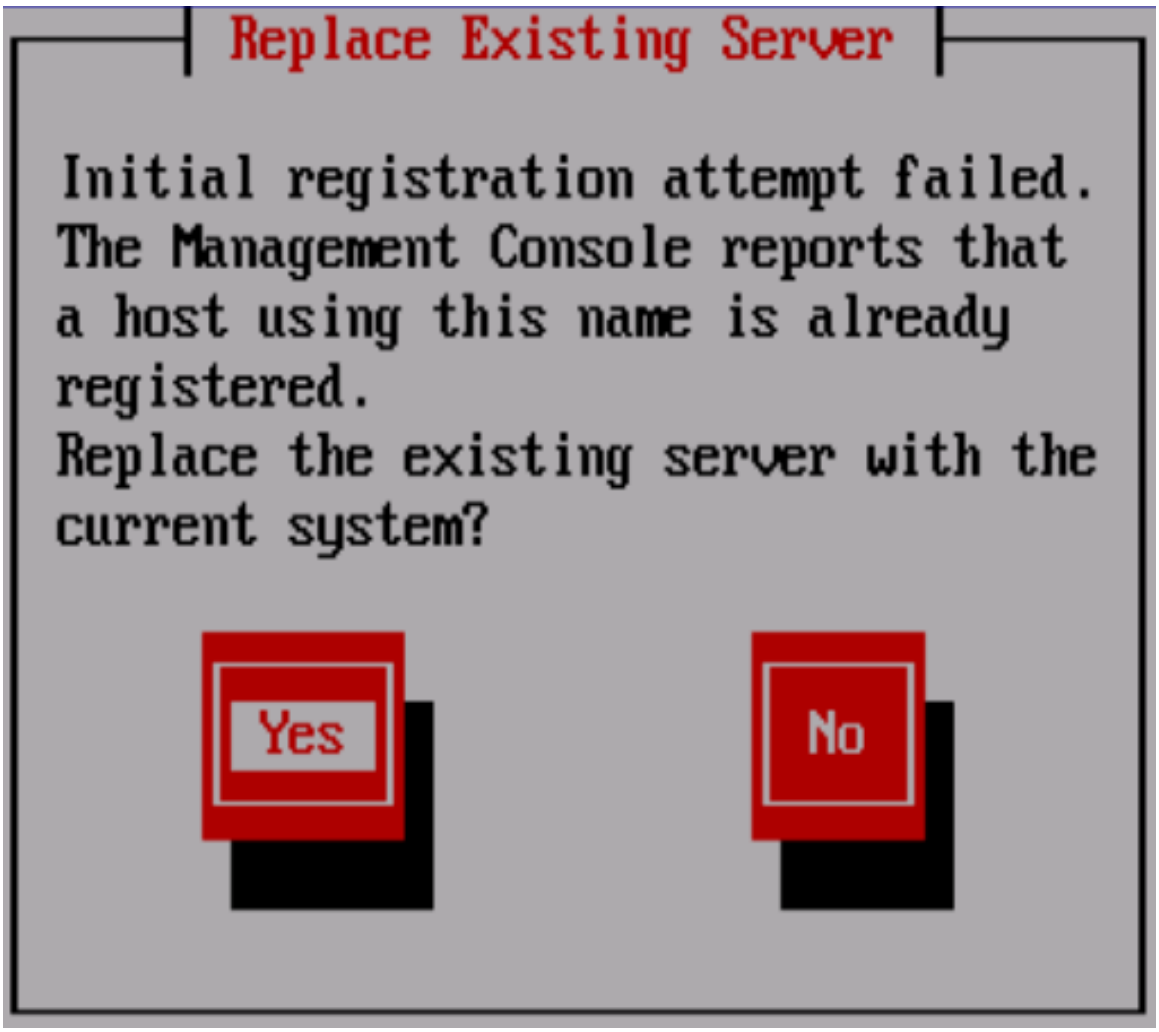
Replace Existing Server
Initial registration attempt failed.
The Management Console reports that a host using this name is already registered.
Replace the existing server with the current system?
Yes
No

Completing the restore on a file serving node

Complete the following steps:

1. Ensure that you have root access to the node. The restore process sets the root password to **`hpinvent`**, the factory default.
2. Verify information about the node you restored:

   ```
   ibrix_server -f [-p] [-M] [-N] -h SERVERNAME
   ```

3. If you disabled NIC monitoring before using the QuickRestore, re-enable the monitor:

   ```
   ibrix_nic -m -h MONITORHOST -A DESTHOST/IFNAME
   ```

   For example:

   ```
   ibrix_nic -m -h titan16 -A titan15/eth2
   ```

4. Configure Insight Remote Support on the node. See "Configuring HP Insight Remote Support on StoreAll systems" (page 35).
5. Run `ibrix_health -l` from the StoreAll management console to verify that no errors are being reported.

Restoring services

When you perform a Quick Restore of a file serving node, the NFS, SMB, FTP, and HTTP export information is not automatically restored to the node. After operations are failed back to the node, the I/O from client systems to the node fails for the NFS, SMB, FTP, and HTTP shares. To avoid this situation, manually restore the NFS, SMB, FTP, and HTTP exports on the node before failing it back.

Restore SMB services. Complete the following steps:

1. If the restored node was previously configured to perform domain authorization, run the following command:

   ```
   ibrix_auth -n DOMAIN_NAME -A AUTH_PROXY_USER_NAME@domain_name [-P
   AUTH_PROXY_PASSWORD] -h HOSTNAME
   ```

   For example:

   ```
   ibrix_auth -n ibq1.mycompany.com -A Administrator@ibq1.mycompany.com
   -P password -h ib5-9
   ```

   If the command fails, check the following:

   - Verify that DNS services are running on the node where you ran the `ibrix_auth` command.
   - Verify that you entered a valid domain name with the full path for the `-n` and `-A` options.

2. Rejoin the likewise database to the Active Directory domain:

   ```
   /opt/likewise/bin/domainjoin-cli join <domain_name> Administrator
   ```

3. Push the original share information from the management console database to the restored node. On the node hosting the active management console, first create a temporary SMB share:

   ```
   ibrix_cifs -a -f FSNAME -s SHARENAME -p SHAREPATH
   ```

   NOTE: You cannot create an SMB share with a name containing an exclamation point (!) or a number sign (#) or both.

   Then delete the temporary SMB share:

```
ibrix_cifs -d -s SHARENAME
```

4. Run the following command to verify that the original share information is on the restored node:

```
ibrix_cifs -i -h SERVERNAME
```

Restore HTTP services. Complete the following steps:

1. Take the appropriate actions:
   - If Active Directory authentication is used, join the restored node to the AD domain manually.
   - If Local user authentication is used, create a temporary local user on the GUI and apply the settings to all servers. This step resynchronizes the local user database. Then remove the temporary user.
2. Run the following command:

```
ibrix_httpconfig -R -h HOSTNAME
```

3. Verify that HTTP services have been restored. Use the GUI or CLI to identify a share served by the restored node and then browse to the share.

All Vhosts and HTTP shares should now be restored on the node.

Restore FTP services. Complete the following steps:

1. Take the appropriate actions:
   - If Active Directory authentication is used, join the restored node to the AD domain manually.
   - If Local user authentication is used, create a temporary local user on the GUI and apply the settings to all servers. This step resynchronizes the local user database. Then remove the temporary user.
2. Run the following command:

```
ibrix_ftpconfig -R -h HOSTNAME
```

3. Verify that HTTP services have been restored. Use the GUI or CLI to identify a share served by the restored node and then browse to the share.

All Vhosts and FTP shares should now be restored on the node.

The ibrix_auth command fails after a restore

If the `ibrix_auth` command fails after a QR restore of a server in a cluster with a message similar to the following:

```
ibrix_auth -n IBRQA1.HP.COM -A administrator@ibrqa1.hp.com -P password
-h hostnameX

Iad error on host hostnameX failed command (<HIDDEN_COMMAND>) status
(1) output: (Joining to AD Domain:IBRQA1.HP.COM With Computer DNS Name:
hostsnameX.ibrqa1.hp.com )
```

Verify that the content of the `/etc/resolve.conf` file is not empty. If the content is empty, copy the contents of the `/etc/resolve.conf` file on another server to the empty `resolve.conf` file.

16 Support and other resources

Contacting HP

For worldwide technical support information, see the HP support website:

http://www.hp.com/support

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Related information

Using the 9320 Storage

- *HP StoreAll Storage File System User Guide*
- *HP StoreAll Storage CLI Reference*
- *HP StoreAll Storage Release Notes*

To find these documents, go to the StoreAll Manuals page:http://www.hp.com/support/StoreAllManuals

Using and maintaining file serving nodes

- *HP ProLiant DL380 G7 Server User Guide*
- *HP ProLiant DL380 G7 Server Maintenance and Service Guide*
- *HP ProLiant DL380 G6 Server User Guide*
- *HP ProLiant DL380 G6 Server Maintenance and Service Guide*

To find these documents, go to the Manuals page (http://www.hp.com/support/manuals) and select **servers** > **ProLiant ml/dl and tc series servers** > **HP ProLiant DL380 G7 Server series** or **HP ProLiant DL380 G6 Server series**.

Using and maintaining the optional dedicated Management Server

- *HP ProLiant DL360 G7 Server User Guide*
- *HP ProLiant DL360 G7 Server Maintenance and Service Guide*
- *HP ProLiant DL360 G6 Server User Guide*
- *HP ProLiant DL360 G6 Server Maintenance and Service Guide*

To find these documents, go to the Manuals page (http://www.hp.com/support/manuals) and select **servers** > **ProLiant ml/dl and tc series servers** >**HP ProLiant DL360 G7 Server series** or **HP ProLiant DL360 G6 Server series**.

Using HP MSA Disk Arrays

- *HP 2000 G2 Modular Smart Array Reference Guide*
- *HP 2000 G2 Modular Smart Array CLI Reference Guide*
- *HP P2000 G3 MSA System CLI Reference Guide*
- Online help for HP Storage Management Utility (SMU) and Command Line Interface (CLI)

To find these documents, go the Manuals page (http://www.hp.com/support/manuals) and select **storage** >**Disk Storage Systems** > **MSA Disk Arrays** >**HP 2000sa G2 Modular Smart Array** or **HP P2000 G3 MSA Array Systems**.

Obtaining spare parts

For the latest spare parts information, go to http://partsurfer.hp.com. Enter your product SKU number to view a list of parts. If you do not have the SKU number, click the **Hierarchy** tab and navigate to your product to view a list of SKUs.

HP websites

For additional information, see the following HP websites:

- http://www.hp.com/go/StoreAll
- http://www.hp.com
- http://www.hp.com/go/storage
- http://www.hp.com/service_locator
- http://www.hp.com/support/manuals
- http://www.hp.com/support/downloads
- http://www.hp.com/storage/whitepapers

Rack stability

Rack stability protects personnel and equipment.

⚠ **WARNING!** To reduce the risk of personal injury or damage to equipment:

- Extend leveling jacks to the floor.
- Ensure that the full weight of the rack rests on the leveling jacks.
- Install stabilizing feet on the rack.
- In multiple-rack installations, fasten racks together securely.
- Extend only one rack component at a time. Racks can become unstable if more than one component is extended.

Product warranties

For information about HP product warranties, see the warranty information website:

http://www.hp.com/go/storagewarranty

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

http://www.hp.com/go/e-updates

After registering, you will receive email notification of product enhancements, new driver versions, firmware updates, and other product resources.

17 Documentation feedback

HP is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (**docsfeedback@hp.com**). Include the document title and part number, version number, or the URL when submitting your feedback.

A Cascading Upgrades

If you are running a StoreAll version earlier than 5.6, do incremental upgrades as described in the following table. If you are running StoreAll 5.6, upgrade to 6.1 before upgrading to 6.3.

If you are upgrading from	Upgrade to	Where to find additional information
StoreAll version 5.4	StoreAll version 5.5	"Upgrading the StoreAll software to the 5.5 release" (page 167)
StoreAll version 5.5	StoreAll version 5.6	"Upgrading the StoreAll software to the 5.6 release" (page 163)
StoreAll version 5.6	StoreAll version 6.1	"Upgrading the StoreAll software to the 6.1 release" (page 154)

Common issue across all upgrades from StoreAll 5.x

If you are upgrading from a StoreAll 5.x release, ensure that the NFS exports option `subtree_check` is the default export option for every NFS export. The `no_subtree_check` option is not compatible with the StoreAll OS software.

To add the `subtree_check` option, perform the following steps:

1. Unexport the NFS exports:

   ```
   ibrix_exportfs -h <HOSTLIST> -p <CLIENT1:PATHNAME1>
   ```

2. Create NFS exports with the `subtree_check` option:

   ```
   ibrix_exportfs -f <FSNAME> -p <CLIENT1:PATHNAME1> -o "subtree_check"
   ```

   NOTE: Multiple options can be specified by using the `-o` parameter by separating each option by a comma (for example, `-o rw,subtree_check`).

3. Complete steps 1 and 2 for every NFS export.
4. Verify that all NFS exports have the `subtree_check` option set:

   ```
   ibrix_exportfs -l
   ```

Upgrading the StoreAll software to the 6.1 release

This section describes how to upgrade to the latest StoreAll software release. The Fusion Manager and all file serving nodes must be upgraded to the new release at the same time.

Upgrades to the StoreAll software 6.1 release are supported for systems currently running StoreAll software 5.6.x and 6.x.

If your system is currently running StoreAll software 5.4.x, first upgrade to 5.5.x, then upgrade to 5.6.x, and then upgrade to 6.1. See "Upgrading the StoreAll software to the 5.5 release" (page 167).

If your system is currently running StoreAll software 5.5.x, upgrade to 5.6.x and then upgrade to 6.1. See "Upgrading the StoreAll software to the 5.6 release" (page 163).

IMPORTANT: If you are upgrading from a StoreAll 5.x release:

- Ensure that the NFS exports option `subtree_check` is the default export option for every NFS export. See "Common issue across all upgrades from StoreAll 5.x" (page 154) for more information.
- Any support tickets collected with the `ibrix_supportticket` command will be deleted during the upgrade. Before upgrading to 6.1.4, download a copy of the archive files (`.tgz`) from the `/admin/platform/diag/supporttickets` directory.

NOTE:

- Verify that the root partition contains adequate free space for the upgrade. Approximately 4 GB is required.
- Be sure to enable password-less access among the cluster nodes before starting the upgrade.
- Do not change the active/passive Fusion Manager configuration during the upgrade.
- Linux StoreAll clients must be upgraded to the 6.x release.

Online upgrades for StoreAll software 6.x to 6.1

Online upgrades are supported only from the StoreAll 6.x release. Upgrades from earlier StoreAll releases must use the appropriate offline upgrade procedure.

When performing an online upgrade, note the following:

- File systems remain mounted and client I/O continues during the upgrade.
- The upgrade process takes approximately 45 minutes, regardless of the number of nodes.
- The total I/O interruption per node IP is four minutes, allowing for a failover time of two minutes and a failback time of two additional minutes.
- Client I/O having a timeout of more than two minutes is supported.

Preparing for the upgrade

To prepare for the upgrade, complete the following steps:

1. Ensure that all nodes are up and running. To determine the status of your cluster nodes, check the dashboard on the GUI or use the `ibrix_health` command.
2. Ensure that High Availability is enabled on each node in the cluster.
3. Verify that `ssh` shared keys have been set up. To do this, run the following command on the node hosting the active instance of the agile Fusion Manager:

   ```
   ssh <server_name>
   ```

   Repeat this command for each node in the cluster and verify that you are not prompted for a password at any time.
4. Ensure that no active tasks are running. Stop any active remote replication, data tiering, or rebalancer tasks running on the cluster. (Use `ibrix_task -l` to list active tasks.) When the upgrade is complete, you can start the tasks again.
5. The 6.1 release requires that nodes hosting the agile management be registered on the cluster network. Run the following command to verify that nodes hosting the agile Fusion Manager have IP addresses on the cluster network:

   ```
   ibrix_fm -f
   ```

   If a node is configured on the user network, see "Node is not registered with the cluster network " (page 22) for a workaround.

Performing the upgrade

The online upgrade is supported only from the StoreAll 6.x to 6.1 release. Complete the following steps:

1. Obtain the latest HP StoreAll 6.1 ISO image from the StoreAll software dropbox. Contact HP Support to register for the release and obtain access to the dropbox.
2. Mount the ISO image and copy the entire directory structure to the `/root/ibrix` directory on the disk running the OS.
3. Change directory to `/root/ibrix` on the disk running the OS and then run `chmod -R 777 *` on the entire directory structure.
4. Run the upgrade script and follow the on-screen directions:

   ```
   ./auto_online_ibrixupgrade
   ```

5. Upgrade Linux StoreAll clients. See "Upgrading Linux StoreAll clients" (page 18).
6. If you received a new license from HP, install it as described in the "Licensing" chapter in this guide.

After the upgrade

Complete these steps:

1. Start any remote replication, rebalancer, or data tiering tasks that were stopped before the upgrade.
2. If your cluster includes G6 servers, check the iLO2 firmware version. The firmware must be at version 2.05 for HA to function properly. If your servers have an earlier version of the iLO2 firmware, download iLO2 version 2.05 using the following URL and copy the firmware update to each G6 server. Follow the installation instructions noted in the URL. This issue does not affect G7 servers.

   http://h20000.www2.hp.com/bizsupport/TechSupport/SoftwareDescription.jsp?lang=en&cc=us&prodTypeId=15351&prodSeriesId=1146658&swItem=MTX-949698a14e114478b9fe126499&prodNameId=1135772&swEnvOID=4103&swLang=8&taskId=135&mode=3
3. A change in the inode format impacts files used for:
   - **Snapshots.** Files used for snapshots must either be created on StoreAll software 6.0 or later, or the pre-6.0 file system containing the files must be upgraded for snapshots. To upgrade a file system, use the `upgrade60.sh` utility. For more information, see "Upgrading pre-6.0 file systems for software snapshots" (page 159).
   - **Data retention.** Files used for data retention (including WORM and auto-commit) must be created on StoreAll software 6.1.1 or later, or the pre-6.1.1 file system containing the files must be upgraded for retention features. To upgrade a file system, use the `ibrix_reten_adm -u -f` *`FSNAME`* command. Additional steps are required before and after you run the `ibrix_reten_adm -u -f` *`FSNAME`* command. For more information, see "Upgrading pre-6.0 file systems for software snapshots" (page 159).

Offline upgrades for StoreAll software 5.6.x or 6.0.x to 6.1

Preparing for the upgrade

To prepare for the upgrade, complete the following steps:

1. Ensure that all nodes are up and running. To determine the status of your cluster nodes, check the dashboard on the GUI or use the `ibrix_health` command.
2. Verify that `ssh` shared keys have been set up. To do this, run the following command on the node hosting the active instance of the agile Fusion Manager:

   `ssh <server_name>`

   Repeat this command for each node in the cluster.
3. Note any custom tuning parameters, such as file system mount options. When the upgrade is complete, you can reapply the parameters.
4. Ensure that no active tasks are running. Stop any active remote replication, data tiering, or rebalancer tasks running on the cluster. (Use `ibrix_task -l` to list active tasks.) When the upgrade is complete, you can start the tasks again.
5. The 6.1 release requires that nodes hosting the agile management be registered on the cluster network. Run the following command to verify that nodes hosting the agile Fusion Manager have IP addresses on the cluster network:

   `ibrix_fm -f`

   If a node is configured on the user network, see "Node is not registered with the cluster network " (page 22) for a workaround.
6. Stop all client I/O to the cluster or file systems. On the Linux client, use `lsof </mountpoint>` to show open files belonging to active processes.

7. Unmount file systems on Linux StoreAll clients:

   ```
   ibrix_umount -f MOUNTPOINT
   ```

8. On all nodes hosting the passive Fusion Manager, place the Fusion Manager into maintenance mode:

   ```
   <ibrixhome>/bin/ibrix_fm -m maintenance -A
   ```

9. On the active Fusion Manager node, disable automated failover on all file serving nodes:

   ```
   <ibrixhome>/bin/ibrix_server -m -U
   ```

10. Run the following command to verify that automated failover is off. In the output, the HA column should display `off`.

    ```
    <ibrixhome>/bin/ibrix_server -l
    ```

11. Stop the SMB, NFS and NDMP services on all nodes. Run the following commands on the node hosting the active Fusion Manager:

    ```
    ibrix_server -s -t cifs -c stop
    ibrix_server -s -t nfs -c stop
    ibrix_server -s -t ndmp -c stop
    ```

    If you are using SMB, verify that all likewise services are down on all file serving nodes:

    ```
    ps -ef | grep likewise
    ```

    Use `kill -9` to stop any likewise services that are still running.

    If you are using NFS, verify that all NFS processes are stopped:

    ```
    ps -ef | grep nfs
    ```

    If necessary, use the following command to stop NFS services:

    ```
    /etc/init.d/nfs stop
    ```

    Use `kill -9` to stop any NFS processes that are still running.

    If necessary, run the following command on all nodes to find any open file handles for the mounted file systems:

    ```
    lsof </mountpoint>
    ```

    Use `kill -9` to stop any processes that still have open file handles on the file systems.

12. Unmount each file system manually:

    ```
    ibrix_umount -f FSNAME
    ```

    Wait up to 15 minutes for the file systems to unmount.

    Troubleshoot any issues with unmounting file systems before proceeding with the upgrade. See "File system unmount issues" (page 23).

Performing the upgrade

This upgrade method is supported only for upgrades from StoreAll software 5.6.x to the 6.1 release. Complete the following steps:

1. Obtain the latest HP StoreAll 6.1 ISO image from the StoreAll software dropbox. Contact HP Support to register for the release and obtain access to the dropbox.

   Mount the ISO image and copy the entire directory structure to the `/root/ibrix` directory on the disk running the OS.

2. Change directory to `/root/ibrix` on the disk running the OS and then run `chmod -R 777 *` on the entire directory structure.
3. Run the following upgrade script:

   ```
   ./auto_ibrixupgrade
   ```

   The upgrade script automatically stops the necessary services and restarts them when the upgrade is complete. The upgrade script installs the Fusion Manager on all file serving nodes. The Fusion Manager is in active mode on the node where the upgrade was run, and is in

passive mode on the other file serving nodes. If the cluster includes a dedicated Management Server, the Fusion Manager is installed in passive mode on that server.

4. Upgrade Linux StoreAll clients. See "Upgrading Linux StoreAll clients" (page 18).
5. If you received a new license from HP, install it as described in the "Licensing" chapter in this guide.

After the upgrade

Complete the following steps:

1. Run the following command to rediscover physical volumes:

   ```
   ibrix_pv -a
   ```

2. Apply any custom tuning parameters, such as mount options.
3. Remount all file systems:

   ```
   ibrix_mount -f <fsname> -m </mountpoint>
   ```

4. Re-enable High Availability if used:

   ```
   ibrix_server -m
   ```

5. Start any remote replication, rebalancer, or data tiering tasks that were stopped before the upgrade.
6. If you are using SMB, set the following parameters to synchronize the SMB software and the Fusion Manager database:

   - `smb signing enabled`
   - `smb signing required`
   - `ignore_writethru`

   Use `ibrix_cifsconfig` to set the parameters, specifying the value appropriate for your cluster (1=enabled, 0=disabled). The following examples set the parameters to the default values for the 6.1 release:

   ```
   ibrix_cifsconfig -t -S "smb_signing_enabled=0,
   smb_signing_required=0"
   ```

   ```
   ibrix_cifsconfig -t -S "ignore_writethru=1"
   ```

   The SMB signing feature specifies whether clients must support SMB signing to access SMB shares. See the *HP StoreAll Storage File System User Guide* for more information about this feature. When`ignore_writethru` is enabled, StoreAll software ignores writethru buffering to improve SMB write performance on some user applications that request it.
7. Mount file systems on Linux StoreAll clients.
8. Because of a change in the inode format, files used for snapshots must either be created on StoreAll 6.0 or later, or the pre-6.0 file system containing the files must be upgraded for snapshots. For more information about upgrading a file system, see "Upgrading pre-6.0 file systems for software snapshots" (page 159).

Upgrading Linux StoreAll clients

Be sure to upgrade the cluster nodes before upgrading Linux StoreAll clients. Complete the following steps on each client:

1. Download the latest HP StoreAll client 6.1 package.
2. Expand the tar file.
3. Run the upgrade script:

   ```
   ./ibrixupgrade -f
   ```

   The upgrade software automatically stops the necessary services and restarts them when the upgrade is complete.

4. Execute the following command to verify the client is running StoreAll software:

```
/etc/init.d/ibrix_client status
IBRIX Filesystem Drivers loaded
IBRIX IAD Server (pid 3208) running...
```

The IAD service should be running, as shown in the previous sample output. If it is not, contact HP Support.

Installing a minor kernel update on Linux clients

The StoreAll client software is upgraded automatically when you install a compatible Linux minor kernel update.

If you are planning to install a minor kernel update, first run the following command to verify that the update is compatible with the StoreAll client software:

```
/usr/local/ibrix/bin/verify_client_update <kernel_update_version>
```

The following example is for a RHEL 4.8 client with kernel version 2.6.9-89.ELsmp:

```
# /usr/local/ibrix/bin/verify_client_update 2.6.9-89.35.1.ELsmp
Kernel update 2.6.9-89.35.1.ELsmp is compatible.
```

If the minor kernel update is compatible, install the update with the vendor RPM and reboot the system. The StoreAll client software is then automatically updated with the new kernel, and StoreAll client services start automatically. Use the `ibrix_version -l -C` command to verify the kernel version on the client.

NOTE: To use the `verify_client` command, the StoreAll client software must be installed.

Upgrading Windows StoreAll clients

Complete the following steps on each client:

1. Remove the old Windows StoreAll client software using the **Add or Remove Programs** utility in the Control Panel.
2. Copy the Windows StoreAll client MSI file for the upgrade to the machine.
3. Launch the Windows Installer and follow the instructions to complete the upgrade.
4. Register the Windows StoreAll client again with the cluster and check the option to **Start Service after Registration**.
5. Check **Administrative Tools** | **Services** to verify that the StoreAll client service is started.
6. Launch the Windows StoreAll client. On the Active Directory Settings tab, click **Update** to retrieve the current Active Directory settings.
7. Mount file systems using the StoreAll Windows client GUI.

NOTE: If you are using Remote Desktop to perform an upgrade, you must log out and log back in to see the drive mounted.

Upgrading pre-6.0 file systems for software snapshots

To support software snapshots, the inode format was changed in the StoreAll 6.0 release. The `upgrade60.sh` utility upgrades a file system created on a pre-6.0 release, enabling software snapshots to be taken on the file system.

The utility can also determine the needed conversions without actually performing the upgrade.

When using the utility, you should be aware of the following:

- The file system must be unmounted.
- Segments marked as `BAD` are not upgraded.
- The upgrade takes place in parallel across all file serving nodes owning segments in the file system, with at least one thread running on each node. For a system with multiple controllers, the utility will run a thread for each controller if possible.

- Files up to 3.8 TB in size can be upgraded. To enable snapshots on larger files, they must be migrated after the upgrade is complete (see "Migrating large files" (page 160).
- In general, the upgrade takes approximately three hours per TB of data. The configuration of the system can affect this number.

Running the utility

Typically, the utility is run as follows to upgrade a file system:

```
upgrade60.sh file system
```

For example, the following command performs a full upgrade on file system `fs1`:

```
upgrade60.sh fs1
```

Progress and status reports

The utility writes log files to the directory `/usr/local/ibrix/log/upgrade60` on each node containing segments from the file system being upgraded. Each node contains the log files for its segments.

Log files are named `<host>_<segment>_<date>_upgrade.log`. For example, the following log file is for segment `ilv2` on host `ib4-2`:

```
ib4-2_ilv2_2012-03-27_11:01_upgrade.log
```

Restarting the utility

If the upgrade is stopped or the system shuts down, you can restart the upgrade utility and it will continue the operation. (To stop an upgrade, press **Ctrl-C** on the command line or send an interrupt signal to the process.)

There should be no adverse effects to the file system; however, certain blocks that were newly allocated by the file system at the time of the interruption will be lost. Running `ibrix_fsck` in corrective mode will recover the blocks.

NOTE: The `upgrade60.sh` utility cannot upgrade segments in an INACTIVE state. If a node is rebooted or shuts down with an unmounted file system, the file system segments owned by that node will be in an INACTIVE state. To move the segments to ACTIVE states, mount the file system with `ibrix_mount`. Then unmount the file system with `ibrix_umount` and resume running `upgrade60.sh`. You can verify segment states with the Linux `lvscan` command.

Migrating large files

The `upgrade60.sh` utility does not upgrade files larger than 3.8 TB. After the upgrade is complete and the file system is mounted, migrate the file to another segment in the file system using the following command:

```
ibmigrate -f filesystem -m 1 -d destination_segment file
```

The following example migrates `file.9` from its current segment to destination segment 2:

```
ibmigrate -f ibfs -m 1 -d 2 /mnt/storeall/test_dir/dir1/file.9
```

After the file is migrated, you can snap the file.

Synopsis

Run the upgrade utility:

```
upgrade60.sh [-v -n] file system
```

The `-n` option lists needed conversions but does not attempt them. The `-v` option provides more information.

Upgrading pre-6.1.1 file systems for data retention features

Data retention was automatically enabled for file systems created with StoreAll 6.1.1 or later. If you want to enable data retention for file systems created with StoreAll 6.0 or earlier, run the `ibrix_reten_adm -u` command, as described in this section.

To enable data retention:

1. If you have a pre-6.0 file system, run the `upgrade60.sh` utility, as described in Section (page 159).
2. Run the following command on a node that has the file system mounted:

   ```
   ibrix_reten_adm -u -f FSNAME
   ```

   In this instance, `FSNAME` is the name of the file system you want to upgrade for data retention features.

   The command enables data retention and unmounts the file system on the node.
3. After the command finishes upgrading the file system, re-mount the file system.
4. Enter the `ibrix_fs` command to set the file system's data retention and autocommit period to the desired values. See the *HP StoreAll Storage CLI Reference Guide* for additional information about the `ibrix_fs` command.

Troubleshooting upgrade issues

If the upgrade does not complete successfully, check the following items. For additional assistance, contact HP Support.

Automatic upgrade

Check the following:

- If the initial execution of `/usr/local/ibrix/setup/upgrade` fails, check `/usr/local/ibrix/setup/upgrade.log` for errors. It is imperative that all servers are up and running the StoreAll software before you execute the upgrade script.
- If the install of the new OS fails, power cycle the node. Try rebooting. If the install does not begin after the reboot, power cycle the machine and select the `upgrade` line from the `grub` boot menu.
- After the upgrade, check `/usr/local/ibrix/setup/logs/postupgrade.log` for errors or warnings.
- If configuration restore fails on any node, look at `/usr/local/ibrix/autocfg/logs/appliance.log` on that node to determine which feature restore failed. Look at the specific feature log file under `/usr/local/ibrix/setup/logs/` for more detailed information.

  To retry the copy of configuration, use the following command:

  ```
  /usr/local/ibrix/autocfg/bin/ibrixapp upgrade -f -s
  ```
- If the install of the new image succeeds, but the configuration restore fails and you need to revert the server to the previous install, run the following command and then reboot the machine. This step causes the server to boot from the old version (the alternate partition).

  ```
  /usr/local/ibrix/setup/boot_info -r
  ```
- If the public network interface is down and inaccessible for any node, power cycle that node.

NOTE: Each node stores its `ibrixupgrade.log` file in `/tmp`.

Manual upgrade

Check the following:

- If the restore script fails, check `/usr/local/ibrix/setup/logs/restore.log` for details.
- If configuration restore fails, look at `/usr/local/ibrix/autocfg/logs/appliance.log` to determine which feature restore failed. Look at the specific feature log file under `/usr/local/ibrix/setup/logs/` for more detailed information.

  To retry the copy of configuration, use the following command:

  ```
  /usr/local/ibrix/autocfg/bin/ibrixapp upgrade -f -s
  ```

Offline upgrade fails because iLO firmware is out of date

If the iLO2 firmware is out of date on a node, the `auto_ibrixupgrade` script will fail. The `/usr/local/ibrix/setup/logs/auto_ibrixupgrade.log` reports the failure and describes how to update the firmware.

After updating the firmware, run the following command on the node to complete the StoreAll software upgrade:

```
/root/ibrix/ibrix/ibrixupgrade -f
```

Node is not registered with the cluster network

Nodes hosting the agile Fusion Manager must be registered with the cluster network. If the `ibrix_fm` command reports that the IP address for a node is on the user network, you will need to reassign the IP address to the cluster network. For example, the following commands report that node `ib51-101`, which is hosting the active Fusion Manager, has an IP address on the user network (192.168.51.101) instead of the cluster network.

```
[root@ib51-101 ibrix]# ibrix_fm -i
FusionServer: ib51-101 (active, quorum is running)
================================================
[root@ib51-101 ibrix]# ibrix_fm -f
NAME      IP ADDRESS
--------  ----------
ib51-101  192.168.51.101
ib51-102  10.10.51.102
```

1. If the node is hosting the active Fusion Manager, as in this example, stop the Fusion Manager on that node:

   ```
   [root@ib51-101 ibrix]# /etc/init.d/ibrix_fusionmanager stop
   Stopping Fusion Manager Daemon                          [  OK  ]
   [root@ib51-101 ibrix]#
   ```

2. On the node now hosting the active Fusion Manager (`ib51-102` in the example), unregister node `ib51-101`:

   ```
   [root@ib51-102 ~]# ibrix_fm -u ib51-101
   Command succeeded!
   ```

3. On the node hosting the active Fusion Manager, register node `ib51-101` and assign the correct IP address:

   ```
   [root@ib51-102 ~]# ibrix_fm -R ib51-101 -I 10.10.51.101
   Command succeeded!
   ```

   NOTE: When registering a Fusion Manager, be sure the hostname specified with `-R` matches the hostname of the server.

The `ibrix_fm` commands now show that node `ib51-101` has the correct IP address and node `ib51-102` is hosting the active Fusion Manager.

```
[root@ib51-102 ~]# ibrix_fm -f
NAME      IP ADDRESS
--------  ----------
```

```
ib51-101  10.10.51.101
ib51-102  10.10.51.102
[root@ib51-102 ~]# ibrix_fm -i
FusionServer: ib51-102 (active, quorum is running)
================================================
```

File system unmount issues

If a file system does not unmount successfully, perform the following steps on all servers:

1. Run the following commands:

   ```
   chkconfig ibrix_server off
   chkconfig ibrix_ndmp off
   chkconfig ibrix_fusionmanager off
   ```

2. Reboot all servers.
3. Run the following commands to move the services back to the `on` state. The commands do not start the services.

   ```
   chkconfig ibrix_server on
   chkconfig ibrix_ndmp on
   chkconfig ibrix_fusionmanager on
   ```

4. Unmount the file systems and continue with the upgrade procedure.

Upgrading the StoreAll software to the 5.6 release

This section describes how to upgrade to the latest StoreAll software release. The management console and all file serving nodes must be upgraded to the new release at the same time.

Upgrades to the StoreAll software 5.6 release are supported for systems currently running StoreAll software 5.5.x. If your system is running an earlier release, first upgrade to the 5.5 release, and then upgrade to 5.6. The upgrade procedure upgrades the operating system to Red Hat Enterprise Linux 5.5.

IMPORTANT:

- Ensure that the NFS exports option `subtree_check` is the default export option for every NFS export. See "Common issue across all upgrades from StoreAll 5.x" (page 154) for more information.
- Do not start new remote replication jobs while a cluster upgrade is in progress. If replication jobs were running before the upgrade started, the jobs will continue to run without problems after the upgrade completes.

The upgrade to StoreAll software 5.6 is supported only as an offline upgrade. Because it requires an upgrade of the kernel, the local disk must be reformatted. Clients will experience a short interruption to administrative and file system access while the system is upgraded.

There are two upgrade procedures available depending on the current installation. If you have a StoreAll software 5.5 system that was installed through the QR procedure, you can use the automatic upgrade procedure. If you used an upgrade procedure to install your StoreAll software 5.5 system, you must use the manual procedure. To determine if your system was installed using the QR procedure, run the `df` command. If you see separate file systems mounted on `/`, `/local`, `/stage`, and `/alt`, your system was quick-restored and you can use the automated upgrade procedure. If you do not see these mount points, proceed with the manual upgrade process.

- **Automatic upgrades.** This process uses separate partitioned space on the local disk to save node-specific configuration information. After each node is upgraded, its configuration is automatically reapplied.
- **Manual upgrades.** Before each server upgrade, this process requires that you back up the node-specific configuration information from the server onto an external device. After the server is upgraded, you will need to copy and restore the node-specific configuration information manually.

The upgrade takes approximately 45 minutes for 9320 systems with a standard configuration.

NOTE: If you are upgrading from a StoreAll 5.x release, any support tickets collected with the `ibrix_supportticket` command will be deleted during the upgrade. Download a copy of the archive files (`.tgz`) from the `/admin/platform/diag/supporttickets` directory.

Automatic upgrades

All file serving nodes and management consoles must be up when you perform the upgrade. If a node or management console is not up, the upgrade script will fail. To determine the status of your cluster nodes, check the dashboard on the GUI or use the `ibrix_health` command.

To upgrade all nodes in the cluster automatically, complete the following steps:

1. Check the dashboard on the management console GUI to verify that all nodes are up.
2. Obtain the latest release image from the HP kiosk at http://www.software.hp.com/kiosk (you will need your HP-provided login credentials).
3. Copy the release `.iso` file onto the current active management console.
4. Run the following command, specifying the location of the local `iso` copy as the argument:

   ```
   /usr/local/ibrix/setup/upgrade <iso>
   ```

   The upgrade script performs all necessary upgrade steps on every server in the cluster and logs progress in the file `/usr/local/ibrix/setup/upgrade.log`. After the script completes, each server will be automatically rebooted and will begin installing the latest software.
5. After the install is complete, the upgrade process automatically restores node-specific configuration information and the cluster should be running the latest software. If an `UPGRADE FAILED` message appears on the active management console, see the specified log file for details.

Manual upgrades

The manual upgrade process requires external storage that will be used to save the cluster configuration. Each server must be able to access this media directly, not through a network, as the network configuration is part of the saved configuration. HP recommends that you use a USB stick or DVD.

NOTE: Be sure to read all instructions before starting the upgrade procedure.

To determine which node is hosting the agile management console configuration, run the `ibrix_fm -i` command.

Preparing for the upgrade

Complete the following steps:

1. Ensure that all nodes are up and running.
2. On the active management console node, disable automated failover on all file serving nodes:

   ```
   <ibrixhome>/bin/ibrix_server -m -U
   ```
3. Run the following command to verify that automated failover is off. In the output, the HA column should display `off`.

   ```
   <ibrixhome>/bin/ibrix_server -l
   ```
4. On the active management console node, stop the NFS and SMB services on all file serving nodes to prevent NFS and SMB clients from timing out.

   ```
   <ibrixhome>/bin/ibrix _server -s -t cifs -c stop
   <ibrixhome>/bin/ibrix _server -s -t nfs -c stop
   ```

   Verify that all likewise services are down on all file serving nodes:

   ```
   ps -ef | grep likewise
   ```

   Use `kill -9` to kill any likewise services that are still running.

5. If file systems are mounted from a Windows StoreAll client, unmount the file systems using the Windows client GUI.
6. Unmount all StoreAll file systems:

   ```
   <ibrixhome>/bin/ibrix _umount -f <fsname>
   ```

Saving the node configuration

Complete the following steps on each node, starting with the node hosting the active management console:

1. Run `/usr/local/ibrix/setup/save_cluster_config`. This script creates a `tgz` file named `<hostname>_cluser_config.tgz`, which contains a backup of the node configuration.
2. Save the `<hostname>_cluser_config.tgz` file, which is located in `/tmp`, to the external storage media.

Performing the upgrade

Complete the following steps on each node:

1. Obtain the latest Quick Restore image from the HP kiosk at http://www.software.hp.com/kiosk (you will need your HP-provided login credentials).
2. Burn the ISO image to a DVD.
3. Insert the Quick Restore DVD into the server DVD-ROM drive.
4. Restart the server to boot from the DVD-ROM.
5. When the StoreAll Network Storage System screen appears, enter **`qr`** to install the StoreAll software on the file serving node.

   The server reboots automatically after the software is installed. Remove the DVD from the DVD-ROM drive.

Restoring the node configuration

Complete the following steps on each node, starting with the previous active management console:

1. Log in to the node. The configuration wizard should pop up. Escape out of the configuration wizard.
2. Attach the external storage media containing the saved node configuration information.
3. Restore the configuration. Run the following restore script and pass in the `tgz` file containing the node's saved configuration information as an argument:

   ```
   /usr/local/ibrix/setup/restore <saved_config.tgz>
   ```

4. Reboot the node.

Completing the upgrade

Complete the following steps:

1. Remount all StoreAll file systems:

   ```
   <ibrixhome>/bin/ibrix_mount -f <fsname> -m </mountpoint>
   ```

2. Remount all previously mounted StoreAll file systems on Windows StoreAll clients using the Windows client GUI.
3. If automated failover was enabled before the upgrade, turn it back on from the node hosting the active management console:

   ```
   <ibrixhome>/bin/ibrix_server -m
   ```

4. Confirm that automated failover is enabled:

   ```
   <ibrixhome>/bin/ibrix_server -l
   ```

   In the output, `HA` should display `on`.

5. From the node hosting the active management console, perform a manual backup of the upgraded configuration:

   `<ibrixhome>/bin/ibrix_fm -B`

6. Verify that all version indicators match for file serving nodes. Run the following command from the active management console:

   `<ibrixhome>/bin/ibrix_version -l`

   If there is a version mismatch, run the `/ibrix/ibrixupgrade -f` script again on the affected node, and then recheck the versions. The installation is successful when all version indicators match. If you followed all instructions and the version indicators do not match, contact HP Support.

7. Verify the health of the cluster:

   `<ibrixhome>/bin/ibrix_health -l`

   The output should show `Passed / on`.

8. For an agile configuration, on all nodes hosting the passive management console, return the management console to passive mode:

   `<ibrixhome>/bin/ibrix_fm -m passive`

9. If you received a new license from HP, install it as described in the "Licensing" chapter in this document.

Troubleshooting upgrade issues

If the upgrade does not complete successfully, check the following items. For additional assistance, contact HP Support.

Automatic upgrade

Check the following:

- If the initial execution of `/usr/local/ibrix/setup/upgrade` fails, check `/usr/local/ibrix/setup/upgrade.log` for errors. It is imperative that all servers are up and running the StoreAll software before you execute the upgrade script.
- If the install of the new OS fails, power cycle the node. Try rebooting. If the install does not begin after the reboot, power cycle the machine and select the `upgrade` line from the `grub` boot menu.
- After the upgrade, check `/usr/local/ibrix/setup/logs/postupgrade.log` for errors or warnings.
- If configuration restore fails on any node, look at `/usr/local/ibrix/autocfg/logs/appliance.log` on that node to determine which feature restore failed. Look at the specific feature log file under `/usr/local/ibrix/setup/logs/` for more detailed information.

  To retry the copy of configuration, use the command appropriate for your server:

  - A file serving node:

    `/usr/local/ibrix/autocfg/bin/ibrixapp upgrade -s`

  - An agile node (a file serving node hosting the agile management console):

    `/usr/local/ibrix/autocfg/bin/ibrixapp upgrade -f -s`

- If the install of the new image succeeds, but the configuration restore fails and you need to revert the server to the previous install, execute `boot_info -r` and then reboot the machine. This step causes the server to boot from the old version (the alternate partition).
- If the public network interface is down and inaccessible for any node, power cycle that node.

Manual upgrade

Check the following:

- If the restore script fails, check `/usr/local/ibrix/setup/logs/restore.log` for details.
- If configuration restore fails, look at `/usr/local/ibrix/autocfg/logs/appliance.log` to determine which feature restore failed. Look at the specific feature log file under `/usr/local/ibrix/setup/logs/` for more detailed information.

  To retry the copy of configuration, use the command appropriate for your server:

  - A file serving node:

    ```
    /usr/local/ibrix/autocfg/bin/ibrixapp upgrade -s
    ```

  - An agile node (a file serving node hosting the agile management console):

    ```
    /usr/local/ibrix/autocfg/bin/ibrixapp upgrade -f -s
    ```

Upgrading the StoreAll software to the 5.5 release

This section describes how to upgrade to the StoreAll software 5.5 release. The management console and all file serving nodes must be upgraded to the new release at the same time.

IMPORTANT:

- Do not start new remote replication jobs while a cluster upgrade is in progress. If replication jobs were running before the upgrade started, the jobs will continue to run without problems after the upgrade completes.
- If you are upgrading from a StoreAll 5.x release, ensure that the NFS exports option `subtree_check` is the default export option for every NFS export. See "Common issue across all upgrades from StoreAll 5.x" (page 154) for more information.

NOTE: If you are upgrading from a StoreAll 5.x release, any support tickets collected with the `ibrix_supportticket` command will be deleted during the upgrade. Download a copy of the archive files (`.tgz`) from the `/admin/platform/diag/supporttickets` directory.

Upgrades can be run either online or offline:

- **Online upgrades.** This procedure upgrades the software while file systems remain mounted. Before upgrading a file serving node, you will need to fail the node over to its backup node, allowing file system access to continue. This procedure cannot be used for major upgrades, but is appropriate for minor and maintenance upgrades.
- **Offline upgrades.** This procedure requires that file systems be unmounted on the node and that services be stopped. (Each file serving node may need to be rebooted if NFS or SMB causes the unmount operation to fail.) You can then perform the upgrade. Clients experience a short interruption to file system access while each file serving node is upgraded.

You can use an automatic or a manual procedure to perform an offline upgrade. Online upgrades must be performed manually.

Automatic upgrades

The automated upgrade procedure is run as an offline upgrade. When each file serving node is upgraded, all file systems are unmounted from the node and services are stopped. Clients will experience a short interruption to file system access while the node is upgraded.

All file serving nodes and management consoles must be up when you perform the upgrade. If a node or management console is not up, the upgrade script will fail and you will need to use a manual upgrade procedure instead. To determine the status of your cluster nodes, check the dashboard on the GUI.

To upgrade all nodes in the cluster automatically, complete the following steps:

1. Check the dashboard on the management console GUI to verify that all nodes are up.
2. Verify that you have an even number of FSNs configured in a couplet pair high availability architecture by running the following command:

   ```
   ibrix_server -l
   ```
3. On the current active management console, move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous StoreAll installation on this node, the installer is in `/root/ibrix`.
4. On the current active management console, expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
5. Change to the installer directory on the active management console, if necessary. Run the following command:

   ```
   ./auto_ibrixupgrade
   ```

   The upgrade script performs all necessary upgrade steps on every server in the cluster and logs progress in the `upgrade.log` file. The log file is located in the installer directory.

Manual upgrades

Upgrade paths

There are two manual upgrade paths: a standard upgrade and an agile upgrade.

- The standard upgrade is used on clusters having a dedicated Management Server machine or blade running the management console software.
- The agile upgrade is used on clusters having an agile management console configuration, where the management console software is installed in an active/passive configuration on two cluster nodes.

To determine whether you have an agile management console configuration, run the `ibrix_fm -i` command. If the output reports the status as `quorum is not configured`, your cluster does not have an agile configuration.

Be sure to use the upgrade procedure corresponding to your management console configuration:

- For standard upgrades, use Page 168.
- For agile upgrades, use Page 172.

Online and offline upgrades

Online and offline upgrade procedures are available for both the standard and agile upgrades:

- **Online upgrades.** This procedure upgrades the software while file systems remain mounted. Before upgrading a file serving node, you will need to fail the node over to its backup node, allowing file system access to continue. This procedure cannot be used for major upgrades, but is appropriate for minor and maintenance upgrades.
- **Offline upgrades.** This procedure requires that you first unmount file systems and stop services. (Each file serving node may need to be rebooted if NFS or SMB causes the unmount operation to fail.) You can then perform the upgrade. Clients will experience a short interruption to file system access while each file serving node is upgraded.

Standard upgrade for clusters with a dedicated Management Server machine or blade

Use these procedures if your cluster has a dedicated Management Server machine or blade hosting the management console software. The StoreAll software 5.4.x to 5.5 upgrade can be performed either online or offline. Future releases may require offline upgrades.

NOTE: Be sure to read all instructions before starting the upgrade procedure.

Standard online upgrade

The management console must be upgraded first. You can then upgrade file serving nodes and StoreAll clients in any order.

Upgrading the management console

Complete the following steps on the Management Server machine or blade:

1. Disable automated failover on all file serving nodes:

   ```
   <ibrixhome>/bin/ibrix_server -m -U
   ```

2. Verify that automated failover is off:

   ```
   <ibrixhome>/bin/ibrix_server -l
   ```

   In the output, the `HA` column should display `off`.

3. Move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous StoreAll installation on this node, the installer is in `/root/ibrix`.
4. Expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
5. Change to the installer directory if necessary and run the upgrade:

   ```
   ./ibrixupgrade -f
   ```

6. Verify that the management console is operational:

   ```
   /etc/init.d/ibrix_fusionmanager status
   ```

   The `status` command should report that the correct services are running. The output is similar to this:

   ```
   Fusion Manager Daemon (pid 18748) running...
   ```

7. Check `/usr/local/ibrix/log/fusionserver.log` for errors.

Upgrading file serving nodes

After the management console has been upgraded, complete the following steps on each file serving node:

1. From the management console, manually fail over the file serving node:

   ```
   <ibrixhome>/bin/ibrix_server -f -p -h HOSTNAME
   ```

   The node reboots automatically.

2. Move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous StoreAll installation on this node, the installer is in `/root/ibrix`.
3. Expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
4. Change to the installer directory if necessary and execute the following command:

   ```
   ./ibrixupgrade -f
   ```

   The upgrade automatically stops services and restarts them when the process is complete.

5. When the upgrade is complete, verify that the StoreAll software services are running on the node:

   ```
   /etc/init.d/ibrix_server status
   ```

   The output is similar to the following. If the IAD service is not running on your system, contact HP Support.

   ```
   IBRIX Filesystem Drivers loaded
   ibrcud is running.. pid 23325
   IBRIX IAD Server (pid 23368) running...
   ```

6. Verify that the `ibrix` and `ipfs` services are running:

   ```
   lsmod|grep ibrix
   ibrix 2323332 0 (unused)
   lsmod|grep ipfs
   ipfs1 102592 0 (unused)
   ```

   If either `grep` command returns empty, contact HP Support.
7. From the management console, verify that the new version of StoreAll software FS/IAS is installed on the file serving node:

   ```
   <ibrixhome>/bin/ibrix_version -l -S
   ```
8. If the upgrade was successful, failback the file serving node:

   ```
   <ibrixhome>/bin/ibrix_server -f -U -h HOSTNAME
   ```
9. Repeat steps 1 through 8 for each file serving node in the cluster.

After all file serving nodes have been upgraded and failed back, complete the upgrade.

Completing the upgrade

1. From the management console, turn automated failover back on:

   ```
   <ibrixhome>/bin/ibrix_server -m
   ```
2. Confirm that automated failover is enabled:

   ```
   <ibrixhome>/bin/ibrix_server -l
   ```

   In the output, HA displays `on`.
3. Verify that all version indicators match for file serving nodes. Run the following command from the management console:

   ```
   <ibrixhome>/bin/ibrix_version -l
   ```

   If there is a version mismatch, run the `/ibrix/ibrixupgrade -f` script again on the affected node, and then recheck the versions. The installation is successful when all version indicators match. If you followed all instructions and the version indicators do not match, contact HP Support.
4. Propagate a new segment map for the cluster:

   ```
   <ibrixhome>/bin/ibrix_dbck -I -f FSNAME
   ```
5. Verify the health of the cluster:

   ```
   <ibrixhome>/bin/ibrix_health -l
   ```

   The output should specify `Passed / on`.

Standard offline upgrade

This upgrade procedure is appropriate for major upgrades. The management console must be upgraded first. You can then upgrade file serving nodes in any order.

Preparing for the upgrade

1. From the management console, disable automated failover on all file serving nodes:

   ```
   <ibrixhome>/bin/ibrix_server -m -U
   ```
2. From the management console, verify that automated failover is off. In the output, the `HA` column should display `off`.

   ```
   <ibrixhome>/bin/ibrix_server -l
   ```
3. Stop the NFS and SMB services on all file serving nodes to prevent NFS and SMB clients from timing out:

   ```
   <ibrixhome>/bin/ibrix_server -s -t cifs -c stop
   <ibrixhome>/bin/ibrix_server -s -t nfs -c stop
   ```

Verify that all likewise services are down on all file serving nodes:

```
ps -ef | grep likewise
```

Use `kill -9` to kill any likewise services that are still running.

4. From the management console, unmount all StoreAll file systems:

```
<ibrixhome>/bin/ibrix_umount -f <fsname>
```

Upgrading the management console

Complete the following steps on the management console:

1. Force a backup of the configuration:

   ```
   <ibrixhome>/bin/ibrix_fm -B
   ```

   The output is stored at `/usr/local/ibrix/tmp/fmbackup.zip`. Be sure to save this file in a location outside of the cluster.
2. Move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous StoreAll installation on this node, the installer is in `/root/ibrix`.
3. Expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
4. Change to the installer directory if necessary and execute the following command:

   ```
   ./ibrixupgrade -f
   ```
5. Verify that the management console started successfully:

   ```
   /etc/init.d/ibrix_fusionmanager status
   ```

   The status command confirms whether the correct services are running. Output is similar to the following:

   ```
   Fusion Manager Daemon (pid 18748) running...
   ```
6. Check `/usr/local/ibrix/log/fusionserver.log` for errors.

Upgrading the file serving nodes

After the management console has been upgraded, complete the following steps on each file serving node:

1. Move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous StoreAll installation on this node, the installer is in `/root/ibrix`.
2. Expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
3. Change to the installer directory if necessary and execute the following command:

   ```
   ./ibrixupgrade -f
   ```

   The upgrade automatically stops services and restarts them when the process completes.
4. When the upgrade is complete, verify that the StoreAll software services are running on the node:

   ```
   /etc/init.d/ibrix_server status
   ```

   The output should be similar to the following example. If the IAD service is not running on your system, contact HP Support.

   ```
   IBRIX Filesystem Drivers loaded
   ibrcud is running.. pid 23325
   IBRIX IAD Server (pid 23368) running...
   ```
5. Execute the following commands to verify that the `ibrix` and `ipfs` services are running:

   ```
   lsmod|grep ibrix
   ```

```
ibrix 2323332 0 (unused)
lsmod|grep ipfs
ipfs1 102592 0 (unused)
```

If either `grep` command returns empty, contact HP Support.

6. From the management console, verify that the new version of StoreAll software FS/IAS has been installed on the file serving nodes:

   `<ibrixhome>/bin/ibrix_version -l -S`

Completing the upgrade

1. Remount all file systems:

   `<ibrixhome>/bin/ibrix_mount -f <fsname> -m </mountpoint>`

2. From the management console, turn automated failover back on:

   `<ibrixhome>/bin/ibrix_server -m`

3. Confirm that automated failover is enabled:

   `<ibrixhome>/bin/ibrix_server -l`

   In the output, `HA` displays `on`.

4. From the management console, perform a manual backup of the upgraded configuration:

   `<ibrixhome>/bin/ibrix_fm -B`

5. Verify that all version indicators match for file serving nodes. Run the following command from the management console:

   `<ibrixhome>/bin/ibrix_version -l`

   If there is a version mismatch, run the `/ibrix/ibrixupgrade -f` script again on the affected node, and then recheck the versions. The installation is successful when all version indicators match. If you followed all instructions and the version indicators do not match, contact HP Support.

6. Verify the health of the cluster:

   `<ibrixhome>/bin/ibrix_health -l`

   The output should show `Passed / on`.

Agile upgrade for clusters with an agile management console configuration

Use these procedures if your cluster has an agile management console configuration. The StoreAll software 5.4.x to 5.5 upgrade can be performed either online or offline. Future releases may require offline upgrades.

NOTE: Be sure to read all instructions before starting the upgrade procedure.

Agile online upgrade

Perform the agile online upgrade in the following order:

- File serving node hosting the active management console
- File serving node hosting the passive management console
- Remaining file serving nodes and StoreAll clients

Upgrading the file serving nodes hosting the management console

Complete the following steps:

1. On the node hosting the active management console, force a backup of the management console configuration:

   `<ibrixhome>/bin/ibrix_fm -B`

   The output is stored at `/usr/local/ibrix/tmp/fmbackup.zip`. Be sure to save this file in a location outside of the cluster.
2. On the active management console node, disable automated failover on all file serving nodes:

   `<ibrixhome>/bin/ibrix_server -m -U`
3. Verify that automated failover is off:

   `<ibrixhome>/bin/ibrix_server -l`

   In the output, the `HA` column should display `off`.
4. On the node hosting the active management console, place the management console into maintenance mode. This step fails over the active management console role to the node currently hosting the passive agile management console.

   `<ibrixhome>/bin/ibrix_fm -m maintenance -A`
5. Wait approximately 60 seconds for the failover to complete, and then run the following command on the node that was the target for the failover:

   `<ibrixhome>/bin/ibrix_fm -i`

   The command should report that the agile management console is now `Active` on this node.
6. From the node on which you failed over the active management console in step 4, change the status of the management console from `maintenance` to `passive`:

   `<ibrixhome>/bin/ibrix_fm -m passive`
7. On the node hosting the active management console, manually fail over the node now hosting the passive management console:

   `<ibrixhome>/bin/ibrix_server -f -p -h HOSTNAME`

   Wait a few minutes for the node to reboot and then run the following command to verify that the failover was successful. The output should report `Up, FailedOver`.

   `<ibrixhome>/bin/ibrix_server -l`
8. On the node hosting the active management console, place the management console into maintenance mode:

   `<ibrixhome>/bin/ibrix_fm -m maintenance -A`

   This step fails back the active management console role to the node currently hosting the passive agile management console (the node that originally was active).
9. Wait approximately 90 seconds for the failover to complete, and then run the following command on the node that was the target for the failover:

   `<ibrixhome>/bin/ibrix_fm -i`

   The command should report that the agile management console is now Active on this node.
10. On the node with the active agile management console, move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous StoreAll installation on this node, the installer is in
    `/root/ibrix`.
11. On the node with the active agile management console, expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in
    `/root`, the installer is in `/root/ibrix`.
12. Change to the installer directory if necessary and run the upgrade:

    `./ibrixupgrade -f`

The installer upgrades both the management console software and the file serving node software on this node.

13. Verify the status of the management console:

    ```
    /etc/init.d/ibrix_fusionmanager status
    ```

    The status command confirms whether the correct services are running. Output will be similar to the following:

    ```
    Fusion Manager Daemon (pid 18748) running...
    ```

    Also run the following command, which should report that the console is Active:

    ```
    <ibrixhome>/bin/ibrix_fm -i
    ```

14. Check `/usr/local/ibrix/log/fusionserver.log` for errors.
15. If the upgrade was successful, failback the file serving node. Run the following command on the node with the active agile management console:

    ```
    <ibrixhome>/bin/ibrix_server -f -U -h HOSTNAME
    ```

16. From the node on which you failed back the active management console in step 8, change the status of the management console from maintenance to passive:

    ```
    <ibrixhome>/bin/ibrix_fm -m passive
    ```

17. If the node with the passive management console is also a file serving node, manually fail over the node from the active management console:

    ```
    <ibrixhome>/bin/ibrix_server -f -p -h HOSTNAME
    ```

    Wait a few minutes for the node to reboot, and then run the following command to verify that the failover was successful. The output should report `Up, FailedOver`.

    ```
    <ibrixhome>/bin/ibrix_server -l
    ```

18. On the node with the passive agile management console, move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous StoreAll installation on this node, the installer is in `/root/ibrix`.
19. On the node hosting the passive agile management console, expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
20. Change to the installer directory if necessary and run the upgrade:

    ```
    ./ibrixupgrade -f
    ```

    The installer upgrades both the management console software and the file serving node software on the node.

21. Verify the status of the management console:

    ```
    /etc/init.d/ibrix_fusionmanager status
    ```

    The status command confirms whether the correct services are running. Output will be similar to the following:

    ```
    Fusion Manager Daemon (pid 18748) running...
    ```

    Also run the following command, which should report that the console is passive:

    ```
    <ibrixhome>/bin/ibrix_fm -i
    ```

22. Check `/usr/local/ibrix/log/fusionserver.log` for errors.
23. If the upgrade was successful, fail back the node. Run the following command on the node with the active agile management console:

    ```
    <ibrixhome>/bin/ibrix_server -f -U -h HOSTNAME
    ```

24. Verify that the agile management console software and the file serving node software are now upgraded on the two nodes hosting the agile management console:

    `<ibrixhome>/bin/ibrix_version -l -S`

    Following is some sample output:

```
Fusion Manager version: 5.5.XXX
  ==============================
  Segment Servers
  ===============
  HOST_NAME  FILE_SYSTEM             IAD/IAS  IAD/FS   OS         KERNEL_VERSION  ARCH
  ---------  ----------------------  -------  -------  ---------  --------------  ----
  ib50-86    5.5.205(9000_5_5)  5.5.XXX  5.5.XXX  GNU/Linux  2.6.18-128.el5  x86_64
  ib50-87    5.5.205(9000_5_5)  5.5.XXX  5.5.XXX  GNU/Linux  2.6.18-128.el5  x86_64
```

You can now upgrade any remaining file serving nodes.

Upgrading remaining file serving nodes

Complete the following steps on each file serving node:

1. Manually fail over the file serving node:

   `<ibrixhome>/bin/ibrix_server -f -p -h HOSTNAME`

   The node will be rebooted automatically.
2. Move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous StoreAll installation on this node, the installer is in `/root/ibrix`.
3. Expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in */root*, the installer is in */root/ibrix*.
4. Change to the installer directory if necessary and execute the following command:

   `./ibrixupgrade -f`

   The upgrade automatically stops services and restarts them when the process is complete.
5. When the upgrade is complete, verify that the StoreAll software services are running on the node:

   `/etc/init.d/ibrix_server status`

   The output will be similar to the following. If the IAD service is not running on your system, contact HP Support.

   ```
   IBRIX Filesystem Drivers loaded
   ibrcud is running.. pid 23325
   IBRIX IAD Server (pid 23368) running...
   ```
6. Verify that the `ibrix` and `ipfs` services are running:

   ```
   lsmod|grep ibrix
   ibrix 2323332 0 (unused)
   lsmod|grep ipfs
   ipfs1 102592 0 (unused)
   ```

   If either `grep` command returns empty, contact HP Support.
7. From the management console, verify that the new version of StoreAll software FS/IAS has been installed on the file serving node:

   `<ibrixhome>/bin/ibrix_version -l -S`
8. If the upgrade was successful, failback the file serving node:

   `<ibrixhome>/bin/ibrix_server -f -U -h HOSTNAME`
9. Repeat steps 1 through 8 for each remaining file serving node in the cluster.

After all file serving nodes have been upgraded and failed back, complete the upgrade.

Completing the upgrade

1. From the node hosting the active management console, turn automated failover back on:

   ```
   <ibrixhome>/bin/ibrix_server -m
   ```

2. Confirm that automated failover is enabled:

   ```
   <ibrixhome>/bin/ibrix_server -l
   ```

   In the output, the HA column should display `on`.

3. Verify that all version indicators match for file serving nodes. Run the following command from the active management console:

   ```
   <ibrixhome>/bin/ibrix_version -l
   ```

   If there is a version mismatch, run the `/ibrix/ibrixupgrade -f` script again on the affected node, and then recheck the versions. The installation is successful when all version indicators match. If you followed all instructions and the version indicators do not match, contact HP Support.

4. Propagate a new segment map for the cluster:

   ```
   <ibrixhome>/bin/ibrix_dbck -I -f FSNAME
   ```

5. Verify the health of the cluster:

   ```
   <ibrixhome>/bin/ibrix_health -l
   ```

   The output should specify `Passed / on`.

Agile offline upgrade

This upgrade procedure is appropriate for major upgrades. Perform the agile offline upgrade in the following order:

- File serving node hosting the active management console
- File serving node hosting the passive management console
- Remaining file serving nodes

NOTE: To determine which node is hosting the active management console, run the following command:

```
<ibrixhome>/bin/ibrix_fm -i
```

Preparing for the upgrade

1. On the active management console node, disable automated failover on all file serving nodes:

   ```
   <ibrixhome>/bin/ibrix_server -m -U
   ```

2. Verify that automated failover is off. In the output, the `HA` column should display `off`.

   ```
   <ibrixhome>/bin/ibrix_server -l
   ```

3. On the active management console node, stop the NFS and SMB services on all file serving nodes to prevent NFS and SMB clients from timing out.

   ```
   <ibrixhome>/bin/ibrix_server -s -t cifs -c stop
   <ibrixhome>/bin/ibrix_server -s -t nfs -c stop
   ```

   Verify that all likewise services are down on all file serving nodes:

   ```
   ps -ef | grep likewise
   ```

   Use `kill -9` to kill any likewise services that are still running.

4. Unmount all StoreAll file systems:

   ```
   <ibrixhome>/bin/ibrix_umount -f <fsname>
   ```

Upgrading the file serving nodes hosting the management console

Complete the following steps:

1. On the node hosting the active management console, force a backup of the management console configuration:

   ```
   <ibrixhome>/bin/ibrix_fm -B
   ```

   The output is stored at `/usr/local/ibrix/tmp/fmbackup.zip`. Be sure to save this file in a location outside of the cluster.
2. On the node hosting the passive management console, place the management console into maintenance mode:

   ```
   <ibrixhome>/bin/ibrix_fm -m maintenance -A
   ```
3. On the active management console node, move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous StoreAll installation on this node, the installer is in `/root/ibrix`.
4. On the active management console node, expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
5. Change to the installer directory if necessary and run the upgrade:

   ```
   ./ibrixupgrade -f
   ```

   The installer upgrades both the management console software and the file serving node software on this node.
6. Verify the status of the management console:

   ```
   /etc/init.d/ibrix_fusionmanager status
   ```

   The status command confirms whether the correct services are running. Output will be similar to the following:

   ```
   Fusion Manager Daemon (pid 18748) running...
   ```
7. Check `/usr/local/ibrix/log/fusionserver.log` for errors.
8. Upgrade the remaining management console node. Move the `ibrix` directory used in the previous release to `ibrix.old`. Then expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
9. Change to the installer directory if necessary and run the upgrade:

   ```
   ./ibrixupgrade -f
   ```

   The installer upgrades both the management console software and the file serving node software on the node.
10. On the node that was just upgraded and has its management console in maintenance mode, move the management console back to passive mode:

    ```
    <ibrixhome>/bin/ibrix_fm -m passive
    ```

    The node now resumes its normal backup operation for the active management console.

Upgrading remaining file serving nodes

Complete the following steps on the remaining file serving nodes:

1. Move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous StoreAll installation on this node, the installer is in `/root/ibrix`.
2. Expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
3. Change to the installer directory if necessary and execute the following command:

   ```
   ./ibrixupgrade -f
   ```

The upgrade automatically stops services and restarts them when the process is complete.

4. When the upgrade is complete, verify that the StoreAll software services are running on the node:

   ```
   /etc/init.d/ibrix_server status
   ```

   The output should be similar to the following example. If the IAD service is not running on your system, contact HP Support.

   ```
   IBRIX Filesystem Drivers loaded
   ibrcud is running.. pid 23325
   IBRIX IAD Server (pid 23368) running...
   ```

5. Execute the following commands to verify that the `ibrix` and `ipfs` services are running:

   ```
   lsmod|grep ibrix
   ibrix 2323332 0 (unused)
   lsmod|grep ipfs
   ipfs1 102592 0 (unused)
   ```

   If either `grep` command returns empty, contact HP Support.

6. From the active management console node, verify that the new version of StoreAll software FS/IAS is installed on the file serving nodes:

   ```
   <ibrixhome>/bin/ibrix_version -l -S
   ```

Completing the upgrade

1. Remount the StoreAll file systems:

   ```
   <ibrixhome>/bin/ibrix_mount -f <fsname> -m </mountpoint>
   ```

2. From the node hosting the active management console, turn automated failover back on:

   ```
   <ibrixhome>/bin/ibrix_server -m
   ```

3. Confirm that automated failover is enabled:

   ```
   <ibrixhome>/bin/ibrix_server -l
   ```

   In the output, `HA` should display `on`.

4. From the node hosting the active management console, perform a manual backup of the upgraded configuration:

   ```
   <ibrixhome>/bin/ibrix_fm -B
   ```

5. Verify that all version indicators match for file serving nodes. Run the following command from the active management console:

   ```
   <ibrixhome>/bin/ibrix_version -l
   ```

   If there is a version mismatch, run the `/ibrix/ibrixupgrade -f` script again on the affected node, and then recheck the versions. The installation is successful when all version indicators match. If you followed all instructions and the version indicators do not match, contact HP Support.

6. Verify the health of the cluster:

   ```
   <ibrixhome>/bin/ibrix_health -l
   ```

   The output should show `Passed / on`.

Troubleshooting upgrade issues

Automatic upgrade fails

Check the `upgrade.log` file to determine the source of the failure. (The log file is located in the installer directory.) If it is not possible to perform the automatic upgrade, continue with the manual upgrade procedure.

ibrixupgrade hangs

The installation can hang because the RPM database is corrupted. This is caused by inconsistencies in the Red Hat Package Manager.

Rebuild the RPM database using the following commands and then attempt the installation again. Note that `rm` is followed by a space and then two underscores, and `rpm` is followed by a space and then two dashes:

```
cd /var/lib/rpm
rm __*
rpm --rebuilddb
```

On the management console, `ibrixupgrade` may also hang if the NFS mount points are stale. In this case, clean up the mount points, reboot the management console, and run the upgrade procedure again.

B Component diagrams for 9300 systems

Front view of file serving node

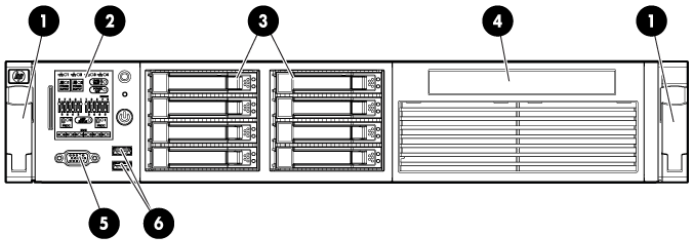

Item	Description
1	Quick-release levers (2)
2	HP Systems Insight Manager display
3	Hard drive bays
4	SATA optical drive bay
5	Video connector
6	USB connectors (2)

Rear view of file serving node

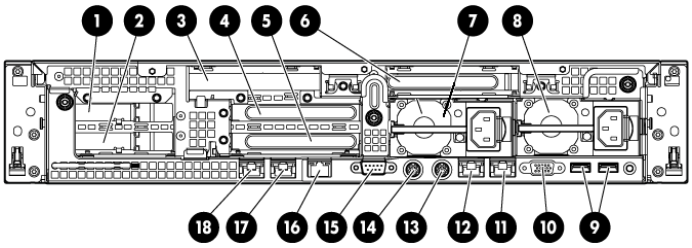

Item	Description
1	PCI slot 5
2	PCI slot 6
3	PCI slot 4
4	PCI slot 2
5	PCI slot 3
6	PCI slot 1
7	Power supply 2 (PS2)
8	Power supply 1 (PS1)
9	USB connectors (2)
10	Video connector

Item	Description
11	NIC 1 connector
12	NIC 2 connector
13	Mouse connector
14	Keyboard connector
15	Serial connector
16	iLO 2 connector
17	NIC 3 connector
18	NIC 4 connector

Server	PCIe card	PCI slot
SATA 1Gb	HP SC08Ge 3Gb SAS Host Bus Adapter	1
	NC364T Quad 1Gb NIC	2
	empty	3
	empty	4
	empty	5
	empty	6
SATA 10Gb	HP SC08Ge 3Gb SAS Host Bus Adapter	1
	empty	2
	empty	3
	NC522SFP dual 10Gb NIC	4
	empty	5
	empty	6
SAS 1Gb	HP SC08Ge 3Gb SAS Host Bus Adapter	1
	NC364T Quad 1Gb NIC	2
	empty	3
	HP SC08Ge 3Gb SAS Host Bus Adapter	4
	empty	5
	empty	6
SAS 10Gb	HP SC08Ge 3Gb SAS Host Bus Adapter	1
	HP SC08Ge 3Gb SAS Host Bus Adapter	2
	empty	3
	NC522SFP dual 10Gb NIC	4
	empty	5
	empty	6

C System component and cabling diagrams for 9320 systems

System component diagrams

Front view of 9300c array controller or 9300cx 3.5" 12-drive enclosure

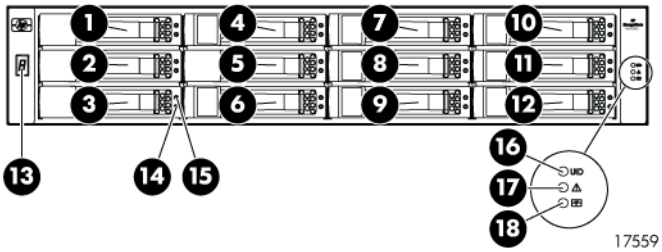

Item	Description
1-12	Disk drive bay numbers
13	Enclosure ID LED
14	Disk drive Online/Activity LED
15	Disk drive Fault/UID LED
16	Unit Identification (UID) LED
17	Fault ID LED
18	Heartbeat ID LED

Rear view of 9300c array controller

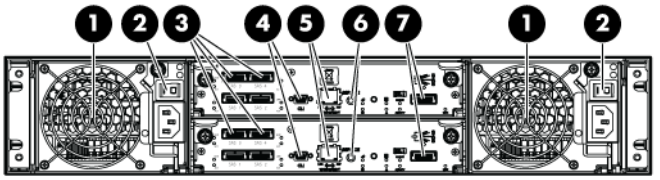

17557

Item	Description
1	Power supplies
2	Power switches
3	Host ports
4	CLI port
5	Network port
6	Service port (used by service personnel only)
7	Expansion port (connects to drive enclosure)

Rear view of 9300cx 3.5" 12-drive enclosure

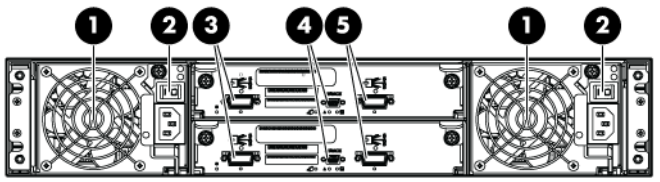

17556

Item	Description
1	Power supplies
2	Power switches
3	SAS In port (connects to the controller enclosure)
4	Service port (used by service personnel only)
5	SAS Out port (connects to another drive enclosure)

Front view of file serving node

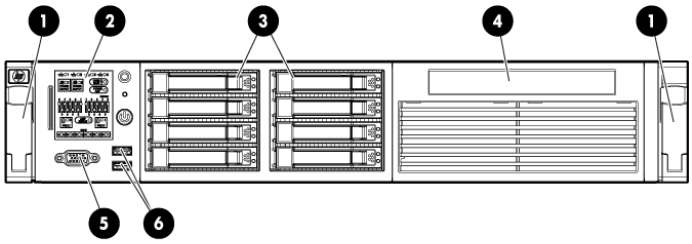

Item	Description
1	Quick-release levers (2)
2	HP Systems Insight Manager display
3	Hard drive bays
4	SATA optical drive bay
5	Video connector
6	USB connectors (2)

Rear view of file serving node

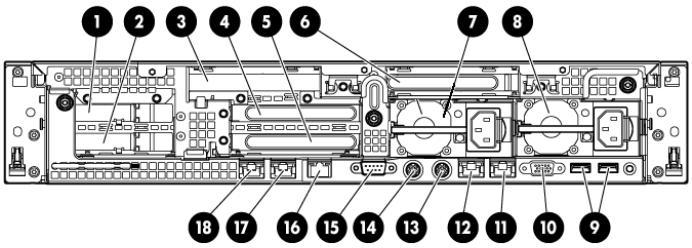

Item	Description
1	PCI slot 5
2	PCI slot 6
3	PCI slot 4
4	PCI slot 2
5	PCI slot 3
6	PCI slot 1
7	Power supply 2 (PS2)
8	Power supply 1 (PS1)
9	USB connectors (2)
10	Video connector
11	NIC 1 connector
12	NIC 2 connector

Item	Description
13	Mouse connector
14	Keyboard connector
15	Serial connector
16	iLO 2 connector
17	NIC 3 connector
18	NIC 4 connector

Server	PCIe card	PCI slot
SATA 1Gb	HP SC08Ge 3Gb SAS Host Bus Adapter	1
	NC364T Quad 1Gb NIC	2
	empty	3
	empty	4
	empty	5
	empty	6
SATA 10Gb	HP SC08Ge 3Gb SAS Host Bus Adapter	1
	empty	2
	empty	3
	NC522SFP dual 10Gb NIC	4
	empty	5
	empty	6
SAS 1Gb	HP SC08Ge 3Gb SAS Host Bus Adapter	1
	NC364T Quad 1Gb NIC	2
	empty	3
	HP SC08Ge 3Gb SAS Host Bus Adapter	4
	empty	5
	empty	6
SAS 10Gb	HP SC08Ge 3Gb SAS Host Bus Adapter	1
	HP SC08Ge 3Gb SAS Host Bus Adapter	2
	empty	3
	NC522SFP dual 10Gb NIC	4
	empty	5
	empty	6

Cabling diagrams

Cluster network cabling diagram

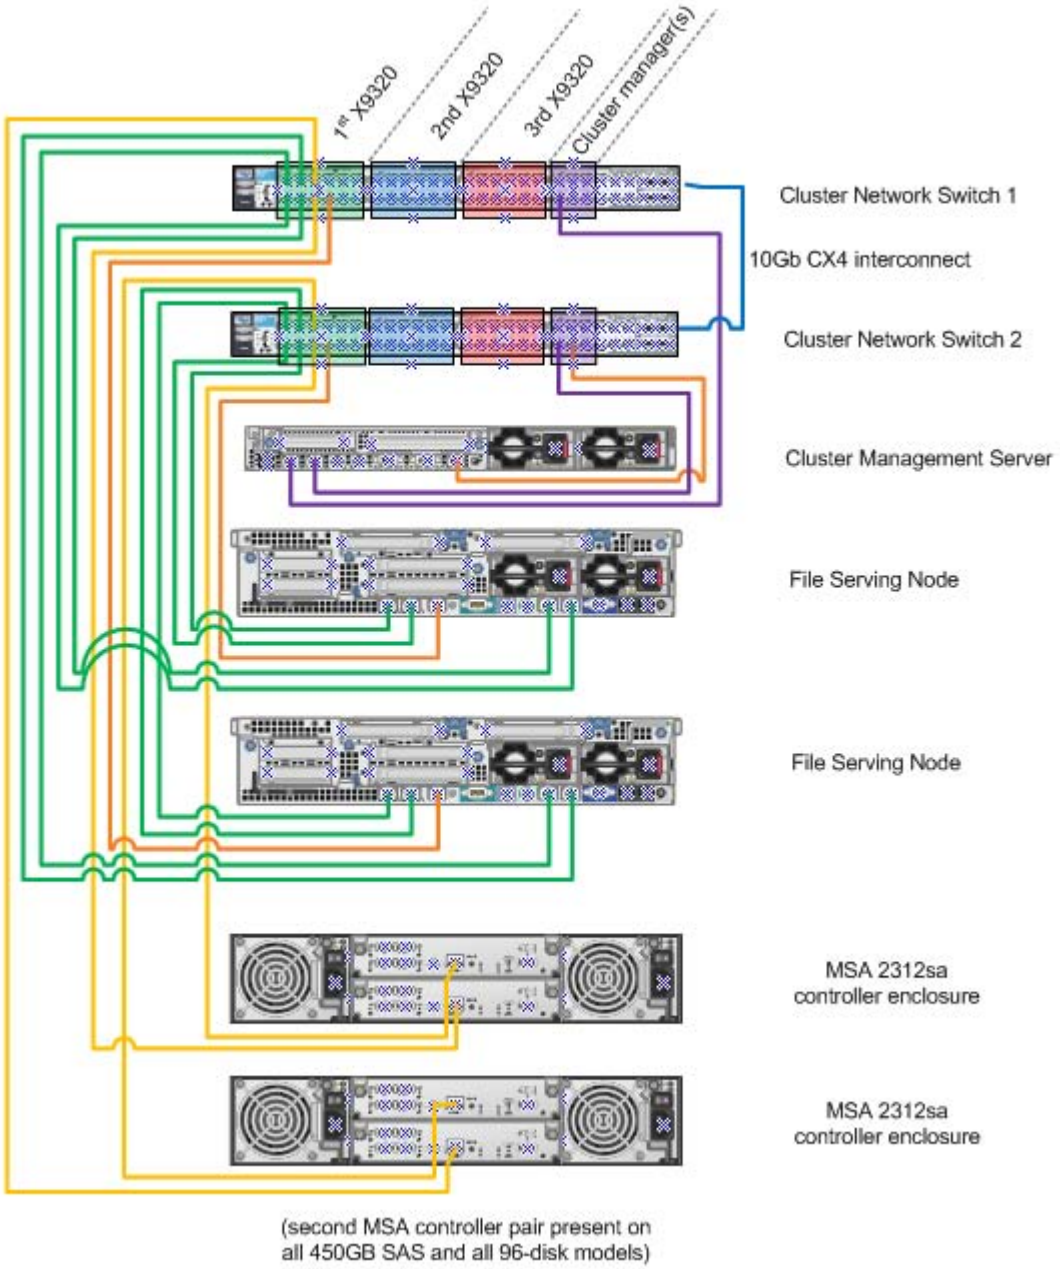

SATA option cabling

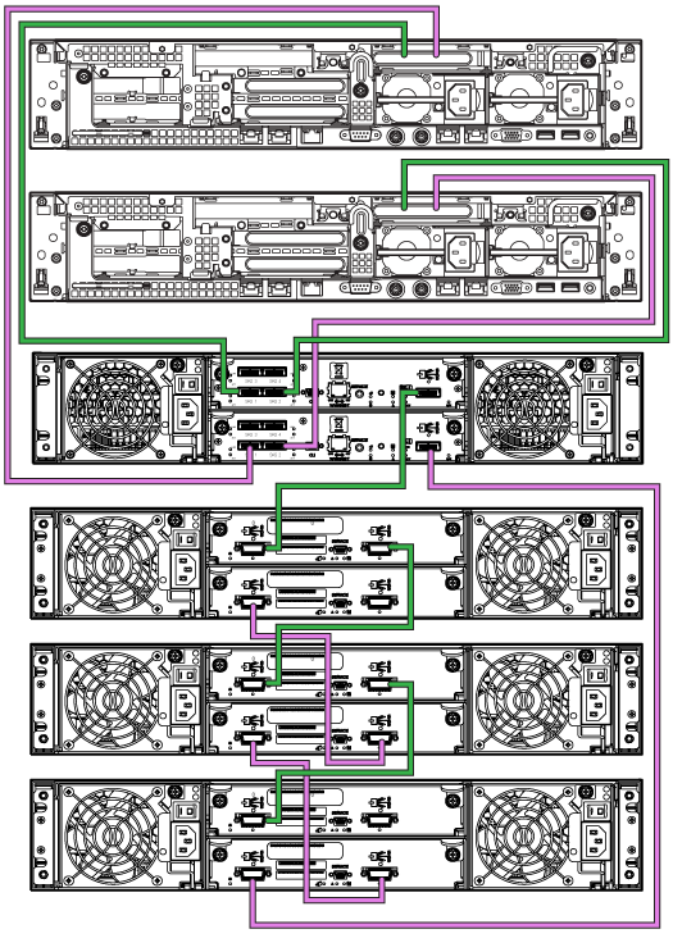

Line	Description
	SAS I/O pathController A
	SAS I/O pathController B

SAS option cabling

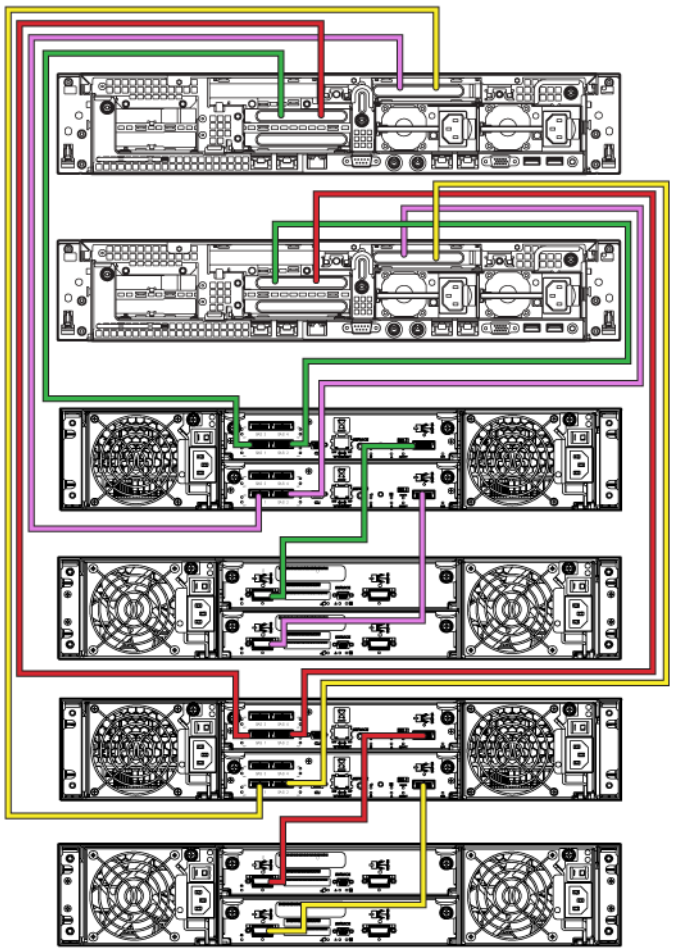

17563

Line	Description
	SAS I/O pathArray 1: Controller A
	SAS I/O pathArray 1: Controller B
	SAS I/O pathArray 2: Controller A
	SAS I/O pathArray 2: Controller B

Drive enclosure cabling

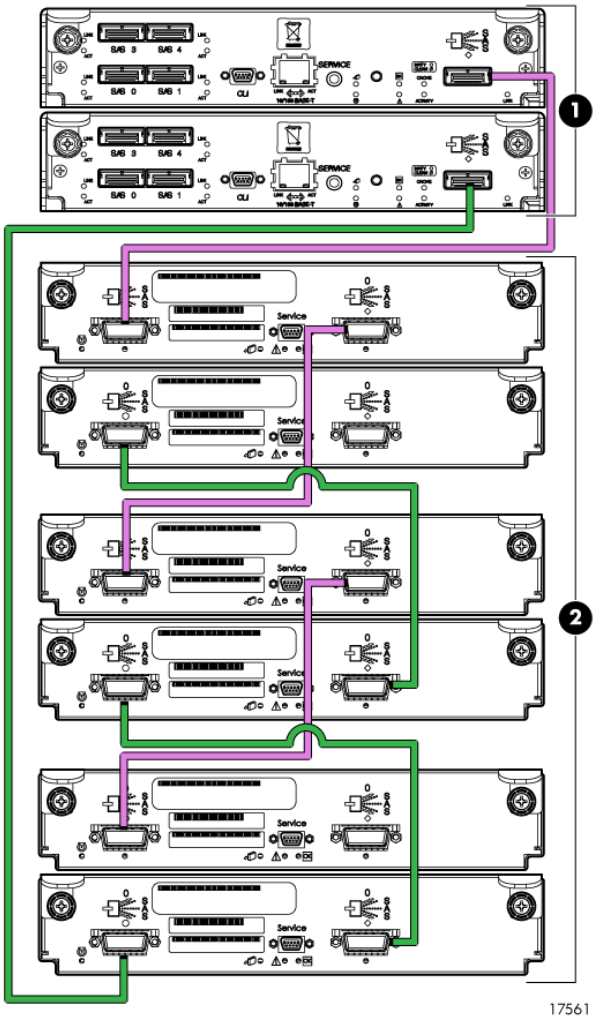

Item	Description
1	SAS controller in 9300c controller enclosure
2	I/O modules in four 9300cx drive enclosures

D Warnings and precautions

Electrostatic discharge information

To prevent damage to the system, be aware of the precautions you need to follow when setting up the system or handling parts. A discharge of static electricity from a finger or other conductor could damage system boards or other static-sensitive devices. This type of damage could reduce the life expectancy of the device.

Preventing electrostatic discharge

To prevent electrostatic damage, observe the following precautions:

- Avoid hand contact by transporting and storing products in static-safe containers.
- Keep electrostatic-sensitive parts in their containers until they arrive at static-free workstations.
- Place parts on a grounded surface before removing them from their containers.
- Avoid touching pins, leads, or circuitry.
- Always be properly grounded when touching a static-sensitive component or assembly.

Grounding methods

There are several methods for grounding. Use one or more of the following methods when handling or installing electrostatic sensitive parts:

- Use a wrist strap connected by a ground cord to a grounded workstation or computer chassis. Wrist straps are flexible straps with a minimum of 1 megohm +/- 10 percent resistance in the ground cords. To provide proper ground, wear the strap snug against the skin.
- Use heel straps, toe straps, or boot straps at standing workstations. Wear the straps on both feet when standing on conductive floors or dissipating floor mats.
- Use conductive field service tools.
- Use a portable field service kit with a folding static-dissipating work mat.

If you do not have any of the suggested equipment for proper grounding, have an HP-authorized reseller install the part.

NOTE: For more information on static electricity or assistance with product installation, contact your HP-authorized reseller.

Equipment symbols

If the following symbols are located on equipment, hazardous conditions could exist.

WARNING!

Any enclosed surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts. To reduce the risk of injury from electrical shock hazards, do not open this enclosure.

WARNING! 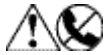

Any RJ-45 receptacle marked with these symbols indicates a network interface connection. To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.

WARNING!

Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. Contact with this surface could result in injury.

WARNING!

Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

WARNING!

Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

Rack warnings and precautions

Ensure that precautions have been taken to provide for rack stability and safety. It is important to follow these precautions providing for rack stability and safety, and to protect both personnel and property. Follow all cautions and warnings included in the installation instructions.

WARNING! To reduce the risk of personal injury or damage to the equipment:

- Observe local occupational safety requirements and guidelines for heavy equipment handling.
- Obtain adequate assistance to lift and stabilize the product during installation or removal.
- Extend the leveling jacks to the floor.
- Rest the full weight of the rack on the leveling jacks.
- Attach stabilizing feet to the rack if it is a single-rack installation.
- Ensure the racks are coupled in multiple-rack installations.
- Fully extend the bottom stabilizers on the equipment. Ensure that the equipment is properly supported/braced when installing options and boards.
- Be careful when sliding rack components with slide rails into the rack. The slide rails could pinch your fingertips.
- Ensure that the rack is adequately stabilized before extending a rack component with slide rails outside the rack. Extend only one component at a time. A rack could become unstable if more than one component is extended for any reason.

WARNING! Verify that the AC power supply branch circuit that provides power to the rack is not overloaded. Overloading AC power to the rack power supply circuit increases the risk of personal injury, fire, or damage to the equipment. The total rack load should not exceed 80 percent of the branch circuit rating. Consult the electrical authority having jurisdiction over your facility wiring and installation requirements.

Device warnings and precautions

WARNING! To reduce the risk of electric shock or damage to the equipment:

- Allow the product to cool before removing covers and touching internal components.
- Do not disable the power cord grounding plug. The grounding plug is an important safety feature.
- Plug the power cord into a grounded (earthed) electrical outlet that is easily accessible at all times.
- Disconnect power from the device by unplugging the power cord from either the electrical outlet or the device.
- Do not use non-conductive tools that could bridge live parts.
- Remove all watches, rings, or loose jewelry when working in hot-plug areas of an energized device.
- Install the device in a controlled access location where only qualified personnel have access to the device.
- Power off the equipment and disconnect power to all AC power cords before removing any access covers for non-hot-pluggable areas.
- Do not replace non-hot-pluggable components while power is applied to the product. Power off the device and then disconnect all AC power cords.
- Do not exceed the level of repair specified in the procedures in the product documentation. All troubleshooting and repair procedures are detailed to allow only subassembly or module-level repair. Because of the complexity of the individual boards and subassemblies, do not attempt to make repairs at the component level or to make modifications to any printed wiring board. Improper repairs can create a safety hazard.

WARNING! To reduce the risk of personal injury or damage to the equipment, the installation of non-hot-pluggable components should be performed only by individuals who are qualified in servicing computer equipment, knowledgeable about the procedures and precautions, and trained to deal with products capable of producing hazardous energy levels.

WARNING! To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

△ **CAUTION:** Protect the installed solution from power fluctuations and temporary interruptions with a regulating Uninterruptible Power Supply (UPS). This device protects the hardware from damage caused by power surges and voltage spikes, and keeps the system in operation during a power failure.

CAUTION: To properly ventilate the system, you must provide at least 7.6 centimeters (3.0 inches) of clearance at the front and back of the device.

CAUTION: Schedule physical configuration changes during periods of low or no activity. If the system is performing rebuilds, RAID migrations, array expansions LUN expansions, or experiencing heavy I/O, avoid physical configuration changes such as adding or replacing hard drives or hot-plugging a controller or any other component. For example, hot-adding or replacing a controller while under heavy I/O could cause a momentary pause, performance decrease, or loss of access to the device while the new controller is starting up. When the controller completes the startup process, full functionality is restored.

CAUTION: Before replacing a hot-pluggable component, ensure that steps have been taken to prevent loss of data.

E Regulatory compliance notices

Regulatory compliance identification numbers

For the purpose of regulatory compliance certifications and identification, this product has been assigned a unique regulatory model number. The regulatory model number can be found on the product nameplate label, along with all required approval markings and information. When requesting compliance information for this product, always refer to this regulatory model number. The regulatory model number is not the marketing name or model number of the product.

Product specific information:

HP _____

Regulatory model number: _____

FCC and CISPR classification: _____

These products contain laser components. See Class 1 laser statement in the Laser compliance notices section.

Federal Communications Commission notice

Part 15 of the Federal Communications Commission (FCC) Rules and Regulations has established Radio Frequency (RF) emission limits to provide an interference-free radio frequency spectrum. Many electronic devices, including computers, generate RF energy incidental to their intended function and are, therefore, covered by these rules. These rules place computers and related peripheral devices into two classes, A and B, depending upon their intended installation. Class A devices are those that may reasonably be expected to be installed in a business or commercial environment. Class B devices are those that may reasonably be expected to be installed in a residential environment (for example, personal computers). The FCC requires devices in both classes to bear a label indicating the interference potential of the device as well as additional operating instructions for the user.

FCC rating label

The FCC rating label on the device shows the classification (A or B) of the equipment. Class B devices have an FCC logo or ID on the label. Class A devices do not have an FCC logo or ID on the label. After you determine the class of the device, refer to the corresponding statement.

Class A equipment

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at personal expense.

Class B equipment

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment

off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit that is different from that to which the receiver is connected.
- Consult the dealer or an experienced radio or television technician for help.

Modification

The FCC requires the user to be notified that any changes or modifications made to this device that are not expressly approved by Hewlett-Packard Company may void the user's authority to operate the equipment.

Cables

When provided, connections to this device must be made with shielded cables with metallic RFI/EMI connector hoods in order to maintain compliance with FCC Rules and Regulations.

Canadian notice (Avis Canadien)

Class A equipment

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la class A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

Class B equipment

This Class B digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la class B respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

European Union notice

This product complies with the following EU directives:

- Low Voltage Directive 2006/95/EC
- EMC Directive 2004/108/EC

Compliance with these directives implies conformity to applicable harmonized European standards (European Norms) which are listed on the EU Declaration of Conformity issued by Hewlett-Packard for this product or product family.

This compliance is indicated by the following conformity marking placed on the product:

CE	This marking is valid for non-Telecom products and EU harmonized Telecom products (e.g., Bluetooth).

Certificates can be obtained from http://www.hp.com/go/certificates.

Hewlett-Packard GmbH, HQ-TRE, Herrenberger Strasse 140, 71034 Boeblingen, Germany

Japanese notices

Japanese VCCI-A notice

この装置は，クラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。　　VCCI－A

Japanese VCCI-B notice

この装置は，クラスB情報技術装置です。この装置は，家庭環境で使用することを目的としていますが，この装置がラジオやテレビジョン受信機に近接して使用されると，受信障害を引き起こすことがあります。
取扱説明書に従って正しい取り扱いをして下さい。　　VCCI－B

Japanese VCCI marking

Japanese power cord statement

製品には、同梱された電源コードをお使い下さい。
同梱された電源コードは、他の製品では使用出来ません。

Please use the attached power cord.
The attached power cord is not allowed to use with other product.

Korean notices

Class A equipment

A급 기기 (업무용 정보통신기기)

이 기기는 업무용으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바라며, 만약 잘못판매 또는 구입하였을 때에는 가정용으로 교환하시기 바랍니다.

Class B equipment

B급 기기 (가정용 정보통신기기)

이 기기는 가정용으로 전자파적합등록을 한 기기로서 주거지역에서는 물론 모든지역에서 사용할 수 있습니다.

Taiwanese notices

BSMI Class A notice

警告使用者:

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Taiwan battery recycle statement

Turkish recycling notice

Türkiye Cumhuriyeti: EEE Yönetmeliğine Uygundur

Vietnamese Information Technology and Communications compliance marking

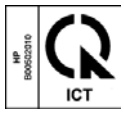

Laser compliance notices

English laser notice

This device may contain a laser that is classified as a Class 1 Laser Product in accordance with U.S. FDA regulations and the IEC 60825-1. The product does not emit hazardous laser radiation.

⚠ **WARNING!** Use of controls or adjustments or performance of procedures other than those specified herein or in the laser product's installation guide may result in hazardous radiation exposure. To reduce the risk of exposure to hazardous radiation:

- Do not try to open the module enclosure. There are no user-serviceable components inside.
- Do not operate controls, make adjustments, or perform procedures to the laser device other than those specified herein.
- Allow only HP Authorized Service technicians to repair the unit.

The Center for Devices and Radiological Health (CDRH) of the U.S. Food and Drug Administration implemented regulations for laser products on August 2, 1976. These regulations apply to laser products manufactured from August 1, 1976. Compliance is mandatory for products marketed in the United States.

Dutch laser notice

WAARSCHUWING: dit apparaat bevat mogelijk een laser die is geclassificeerd als een laserproduct van Klasse 1 overeenkomstig de bepalingen van de Amerikaanse FDA en de richtlijn IEC 60825-1. Dit product geeft geen gevaarlijke laserstraling af.

Als u bedieningselementen gebruikt, instellingen aanpast of procedures uitvoert op een andere manier dan in deze publicatie of in de installatiehandleiding van het laserproduct wordt aangegeven, loopt u het risico te worden blootgesteld aan gevaarlijke straling. Het risico van blootstelling aan gevaarlijke straling beperkt u als volgt:

- Probeer de behuizing van de module niet te openen. U mag zelf geen onderdelen repareren.
- Gebruik voor de laserapparatuur geen andere knoppen of instellingen en voer geen andere aanpassingen of procedures uit dan die in deze handleiding worden beschreven.
- Alleen door HP geautoriseerde technici mogen het apparaat repareren.

French laser notice

AVERTISSEMENT : cet appareil peut être équipé d'un laser classé en tant que Produit laser de classe 1 et conforme à la réglementation de la FDA américaine et à la norme 60825-1 de l'IEC. Ce produit n'émet pas de rayonnement dangereux.

L'utilisation de commandes, de réglages ou de procédures autres que ceux qui sont indiqués ici ou dans le manuel d'installation du produit laser peut exposer l'utilisateur à des rayonnements dangereux. Pour réduire le risque d'exposition à des rayonnements dangereux :

- Ne tentez pas d'ouvrir le boîtier renfermant l'appareil laser. Il ne contient aucune pièce dont la maintenance puisse être effectuée par l'utilisateur.
- Tout contrôle, réglage ou procédure autre que ceux décrits dans ce chapitre ne doivent pas être effectués par l'utilisateur.
- Seuls les Mainteneurs Agréés HP sont habilités à réparer l'appareil laser.

German laser notice

VORSICHT: Dieses Gerät enthält möglicherweise einen Laser, der nach den US-amerikanischen FDA-Bestimmungen und nach IEC 60825-1 als Laserprodukt der Klasse 1 zertifiziert ist. Gesundheitsschädliche Laserstrahlen werden nicht emittiert.

Die Anleitungen in diesem Dokument müssen befolgt werden. Bei Einstellungen oder Durchführung sonstiger Verfahren, die über die Anleitungen in diesem Dokument bzw. im Installationshandbuch des Lasergeräts hinausgehen, kann es zum Austritt gefährlicher Strahlung kommen. Zur Vermeidung der Freisetzung gefährlicher Strahlungen sind die folgenden Punkte zu beachten:

- Versuchen Sie nicht, die Abdeckung des Lasermoduls zu öffnen. Im Inneren befinden sich keine Komponenten, die vom Benutzer gewartet werden können.
- Benutzen Sie das Lasergerät ausschließlich gemäß den Anleitungen und Hinweisen in diesem Dokument.
- Lassen Sie das Gerät nur von einem HP Servicepartner reparieren.

Italian laser notice

AVVERTENZA: AVVERTENZA Questo dispositivo può contenere un laser classificato come prodotto laser di Classe 1 in conformità alle normative US FDA e IEC 60825-1. Questo prodotto non emette radiazioni laser pericolose.

L'eventuale esecuzione di comandi, regolazioni o procedure difformi a quanto specificato nella presente documentazione o nella guida di installazione del prodotto può causare l'esposizione a radiazioni nocive. Per ridurre i rischi di esposizione a radiazioni pericolose, attenersi alle seguenti precauzioni:

- Non cercare di aprire il contenitore del modulo. All'interno non vi sono componenti soggetti a manutenzione da parte dell'utente.
- Non eseguire operazioni di controllo, regolazione o di altro genere su un dispositivo laser ad eccezione di quelle specificate da queste istruzioni.
- Affidare gli interventi di riparazione dell'unità esclusivamente ai tecnici dell'Assistenza autorizzata HP.

Japanese laser notice

警告:本製品には、US FDA規則およびIEC 60825-1に基づくClass 1レーザー製品が含まれている場合があります。本製品は人体に危険なレーザー光は発しません。

本書およびレーザー製品のインストール ガイドに示されている以外の方法で制御、調整、使用した場合、人体に危険な光線にさらされる場合があります。人体に危険な光線にさらされないため、以下の項目を守ってください。

- モジュール エンクロージャを開けないでください。ユーザーが取り扱えるコンポーネントは含まれていません。
- 本書に示されている以外の方法で、レーザー デバイスを制御、調整、使用しないでください。
- HPの正規サービス技術者のみが本ユニットの修理を許可されています。

Spanish laser notice

ADVERTENCIA: Este dispositivo podría contener un láser clasificado como producto de láser de Clase 1 de acuerdo con la normativa de la FDA de EE.UU. e IEC 60825-1. El producto no emite radiaciones láser peligrosas.

El uso de controles, ajustes o manipulaciones distintos de los especificados aquí o en la guía de instalación del producto de láser puede producir una exposición peligrosa a las radiaciones. Para evitar el riesgo de exposición a radiaciones peligrosas:

- No intente abrir la cubierta del módulo. Dentro no hay componentes que el usuario pueda reparar.
- No realice más operaciones de control, ajustes o manipulaciones en el dispositivo láser que los aquí especificados.
- Sólo permita reparar la unidad a los agentes del servicio técnico autorizado HP.

Recycling notices

English recycling notice

Disposal of waste equipment by users in private household in the European Union

This symbol means do not dispose of your product with your other household waste. Instead, you should protect human health and the environment by handing over your waste equipment to a designated collection point for the recycling of waste electrical and electronic equipment. For more information, please contact your household waste disposal service

Bulgarian recycling notice

Изхвърляне на отпадъчно оборудване от потребители в частни домакинства в Европейския съюз

Този символ върху продукта или опаковката му показва, че продуктът не трябва да се изхвърля заедно с другите битови отпадъци. Вместо това, трябва да предпазите човешкото здраве и околната среда, като предадете отпадъчното оборудване в предназначен за събирането му пункт за рециклиране на неизползваемо електрическо и електронно борудване. За допълнителна информация се свържете с фирмата по чистота, чиито услуги използвате.

Czech recycling notice

Likvidace zařízení v domácnostech v Evropské unii

Tento symbol znamená, že nesmíte tento produkt likvidovat spolu s jiným domovním odpadem. Místo toho byste měli chránit lidské zdraví a životní prostředí tím, že jej předáte na k tomu určené sběrné pracoviště, kde se zabývají recyklací elektrického a elektronického vybavení. Pro více informací kontaktujte společnost zabývající se sběrem a svozem domovního odpadu.

Danish recycling notice

Bortskaffelse af brugt udstyr hos brugere i private hjem i EU

Dette symbol betyder, at produktet ikke må bortskaffes sammen med andet husholdningsaffald. Du skal i stedet den menneskelige sundhed og miljøet ved at afl evere dit brugte udstyr på et dertil beregnet indsamlingssted for af brugt, elektrisk og elektronisk udstyr. Kontakt nærmeste renovationsafdeling for yderligere oplysninger.

Dutch recycling notice

Inzameling van afgedankte apparatuur van particuliere huishoudens in de Europese Unie

Dit symbool betekent dat het product niet mag worden gedeponeerd bij het overige huishoudelijke afval. Bescherm de gezondheid en het milieu door afgedankte apparatuur in te leveren bij een hiervoor bestemd inzamelpunt voor recycling van afgedankte elektrische en elektronische apparatuur. Neem voor meer informatie contact op met uw gemeentereinigingsdienst.

Estonian recycling notice

Äravisatavate seadmete likvideerimine Euroopa Liidu eramajapidamistes

See märk näitab, et seadet ei tohi visata olmeprügi hulka. Inimeste tervise ja keskkonna säästmise nimel tuleb äravisatav toode tuua elektriliste ja elektrooniliste seadmete käitlemisega egelevasse kogumispunkti. Küsimuste korral pöörduge kohaliku prügikäitlusettevõtte poole.

Finnish recycling notice

Kotitalousjätteiden hävittäminen Euroopan unionin alueella

Tämä symboli merkitsee, että laitetta ei saa hävittää muiden kotitalousjätteiden mukana. Sen sijaan sinun on suojattava ihmisten terveyttä ja ympäristöä toimittamalla käytöstä poistettu laite sähkö- tai elektroniikkajätteen kierrätyspisteeseen. Lisätietoja saat jätehuoltoyhtiöltä.

French recycling notice

Mise au rebut d'équipement par les utilisateurs privés dans l'Union Européenne

Ce symbole indique que vous ne devez pas jeter votre produit avec les ordures ménagères. Il est de votre responsabilité de protéger la santé et l'environnement et de vous débarrasser de votre équipement en le remettant à une déchetterie effectuant le recyclage des équipements électriques et électroniques. Pour de plus amples informations, prenez contact avec votre service d'élimination des ordures ménagères.

German recycling notice

Entsorgung von Altgeräten von Benutzern in privaten Haushalten in der EU

Dieses Symbol besagt, dass dieses Produkt nicht mit dem Haushaltsmüll entsorgt werden darf. Zum Schutze der Gesundheit und der Umwelt sollten Sie stattdessen Ihre Altgeräte zur Entsorgung einer dafür vorgesehenen Recyclingstelle für elektrische und elektronische Geräte übergeben. Weitere Informationen erhalten Sie von Ihrem Entsorgungsunternehmen für Hausmüll.

Greek recycling notice

Απόρριψη άχρηστου εξοπλισμού από ιδιώτες χρήστες στην Ευρωπαϊκή Ένωση

Αυτό το σύμβολο σημαίνει ότι δεν πρέπει να απορρίψετε το προϊόν με τα λοιπά οικιακά απορρίμματα. Αντίθετα, πρέπει να προστατέψετε την ανθρώπινη υγεία και το περιβάλλον παραδίδοντας τον άχρηστο εξοπλισμό σας σε εξουσιοδοτημένο σημείο συλλογής για την ανακύκλωση άχρηστου ηλεκτρικού και ηλεκτρονικού εξοπλισμού. Για περισσότερες πληροφορίες, επικοινωνήστε με την υπηρεσία απόρριψης απορριμμάτων της περιοχής σας.

Hungarian recycling notice

A hulladék anyagok megsemmisítése az Európai Unió háztartásaiban

Ez a szimbólum azt jelzi, hogy a készüléket nem szabad a háztartási hulladékkal együtt kidobni. Ehelyett a leselejtezett berendezéseknek az elektromos vagy elektronikus hulladék átvételére kijelölt helyen történő beszolgáltatásával megóvja az emberi egészséget és a környezetet.További információt a helyi köztisztasági vállalattól kaphat.

Italian recycling notice

Smaltimento di apparecchiature usate da parte di utenti privati nell'Unione Europea

Questo simbolo avvisa di non smaltire il prodotto con i normali rifi uti domestici. Rispettare la salute umana e l'ambiente conferendo l'apparecchiatura dismessa a un centro di raccolta designato per il riciclo di apparecchiature elettroniche ed elettriche. Per ulteriori informazioni, rivolgersi al servizio per lo smaltimento dei rifi uti domestici.

Latvian recycling notice

Europos Sąjungos namų ūkio vartotojų įrangos atliekų šalinimas

Šis simbolis nurodo, kad gaminio negalima išmesti kartu su kitomis buitinėmis atliekomis. Kad apsaugotumėte žmonių sveikatą ir aplinką, pasenusią nenaudojamą įrangą turite nuvežti į elektrinių ir elektroninių atliekų surinkimo punktą. Daugiau informacijos teiraukitės buitinių atliekų surinkimo tarnybos.

Lithuanian recycling notice

Nolietotu iekārtu iznīcināšanas noteikumi lietotājiem Eiropas Savienības privātajās mājsaimniecībās

Šis simbols norāda, ka ierīci nedrīkst utilizēt kopā ar citiem mājsaimniecības atkritumiem. Jums jārūpējas par cilvēku veselības un vides aizsardzību, nododot lietoto aprīkojumu otrreizējai pārstrādei īpašā lietotu elektrisko un elektronisko ierīču savākšanas punktā. Lai iegūtu plašāku informāciju, lūdzu, sazinieties ar savu mājsaimniecības atkritumu likvidēšanas dienestu.

Polish recycling notice

Utylizacja zużytego sprzętu przez użytkowników w prywatnych gospodarstwach domowych w krajach Unii Europejskiej

Ten symbol oznacza, że nie wolno wyrzucać produktu wraz z innymi domowymi odpadkami. Obowiązkiem użytkownika jest ochrona zdrowa ludzkiego i środowiska przez przekazanie zużytego sprzętu do wyznaczonego punktu zajmującego się recyklingiem odpadów powstałych ze sprzętu elektrycznego i elektronicznego. Więcej informacji można uzyskać od lokalnej firmy zajmującej wywozem nieczystości.

Portuguese recycling notice

Descarte de equipamentos usados por utilizadores domésticos na União Europeia

Este símbolo indica que não deve descartar o seu produto juntamente com os outros lixos domiciliares. Ao invés disso, deve proteger a saúde humana e o meio ambiente levando o seu equipamento para descarte em um ponto de recolha destinado à reciclagem de resíduos de equipamentos eléctricos e electrónicos. Para obter mais informações, contacte o seu serviço de tratamento de resíduos domésticos.

Romanian recycling notice

Casarea echipamentului uzat de către utilizatorii casnici din Uniunea Europeană

Acest simbol înseamnă să nu se arunce produsul cu alte deșeuri menajere. În schimb, trebuie să protejați sănătatea umană și mediul predând echipamentul uzat la un punct de colectare desemnat pentru reciclarea echipamentelor electrice și electronice uzate. Pentru informații suplimentare, vă rugăm să contactați serviciul de eliminare a deșeurilor menajere local.

Slovak recycling notice

Likvidácia vyradených zariadení používateľmi v domácnostiach v Európskej únii

Tento symbol znamená, že tento produkt sa nemá likvidovať s ostatným domovým odpadom. Namiesto toho by ste mali chrániť ľudské zdravie a životné prostredie odovzdaním odpadového zariadenia na zbernom mieste, ktoré je určené na recykláciu odpadových elektrických a elektronických zariadení. Ďalšie informácie získate od spoločnosti zaoberajúcej sa likvidáciou domového odpadu.

Spanish recycling notice

Eliminación de los equipos que ya no se utilizan en entornos domésticos de la Unión Europea

Este símbolo indica que este producto no debe eliminarse con los residuos domésticos. En lugar de ello, debe evitar causar daños a la salud de las personas y al medio ambiente llevando los equipos que no utilice a un punto de recogida designado para el reciclaje de equipos eléctricos y electrónicos que ya no se utilizan. Para obtener más información, póngase en contacto con el servicio de recogida de residuos domésticos.

Swedish recycling notice

Hantering av elektroniskt avfall för hemanvändare inom EU

Den här symbolen innebär att du inte ska kasta din produkt i hushållsavfallet. Värna i stället om natur och miljö genom att lämna in uttjänt utrustning på anvisad insamlingsplats. Allt elektriskt och elektroniskt avfall går sedan vidare till återvinning. Kontakta ditt återvinningsföretag för mer information.

Battery replacement notices

Dutch battery notice

Verklaring betreffende de batterij

WAARSCHUWING: dit apparaat bevat mogelijk een batterij.

- Probeer de batterijen na het verwijderen niet op te laden.
- Stel de batterijen niet bloot aan water of temperaturen boven 60° C.
- De batterijen mogen niet worden beschadigd, gedemonteerd, geplet of doorboord.
- Zorg dat u geen kortsluiting veroorzaakt tussen de externe contactpunten en laat de batterijen niet in aanraking komen met water of vuur.
- Gebruik ter vervanging alleen door HP goedgekeurde batterijen.

Batterijen, accu's en accumulators mogen niet worden gedeponeerd bij het normale huishoudelijke afval. Als u de batterijen/accu's wilt inleveren voor hergebruik of op de juiste manier wilt vernietigen, kunt u gebruik maken van het openbare inzamelingssysteem voor klein chemisch afval of ze terugsturen naar HP of een geautoriseerde HP Business of Service Partner.

Neem contact op met een geautoriseerde leverancier of een Business of Service Partner voor meer informatie over het vervangen of op de juiste manier vernietigen van accu's.

French battery notice

Avis relatif aux piles

AVERTISSEMENT : cet appareil peut contenir des piles.

- N'essayez pas de recharger les piles après les avoir retirées.
- Évitez de les mettre en contact avec de l'eau ou de les soumettre à des températures supérieures à 60°C.
- N'essayez pas de démonter, d'écraser ou de percer les piles.
- N'essayez pas de court-circuiter les bornes de la pile ou de jeter cette dernière dans le feu ou l'eau.
- Remplacez les piles exclusivement par des pièces de rechange HP prévues pour ce produit.

Les piles, modules de batteries et accumulateurs ne doivent pas être jetés avec les déchets ménagers. Pour permettre leur recyclage ou leur élimination, veuillez utiliser les systèmes de collecte publique ou renvoyez-les à HP, à votre Partenaire Agréé HP ou aux agents agréés.

Contactez un Revendeur Agréé ou Mainteneur Agréé pour savoir comment remplacer et jeter vos piles.

German battery notice

Hinweise zu Batterien und Akkus

VORSICHT: Dieses Produkt enthält unter Umständen eine Batterie oder einen Akku.

- Versuchen Sie nicht, Batterien und Akkus außerhalb des Gerätes wieder aufzuladen.
- Schützen Sie Batterien und Akkus vor Feuchtigkeit und Temperaturen über 60°.
- Verwenden Sie Batterien und Akkus nicht missbräuchlich, nehmen Sie sie nicht auseinander und vermeiden Sie mechanische Beschädigungen jeglicher Art.
- Vermeiden Sie Kurzschlüsse, und setzen Sie Batterien und Akkus weder Wasser noch Feuer aus.
- Ersetzen Sie Batterien und Akkus nur durch die von HP vorgesehenen Ersatzteile.

Batterien und Akkus dürfen nicht über den normalen Hausmüll entsorgt werden. Um sie der Wiederverwertung oder dem Sondermüll zuzuführen, nutzen Sie die öffentlichen Sammelstellen, oder setzen Sie sich bezüglich der Entsorgung mit einem HP Partner in Verbindung.

Weitere Informationen zum Austausch von Batterien und Akkus oder zur sachgemäßen Entsorgung erhalten Sie bei Ihrem HP Partner oder Servicepartner.

Italian battery notice

Istruzioni per la batteria

AVVERTENZA: Questo dispositivo può contenere una batteria.

- Non tentare di ricaricare le batterie se rimosse.
- Evitare che le batterie entrino in contatto con l'acqua o siano esposte a temperature superiori a 60° C.
- Non smontare, schiacciare, forare o utilizzare in modo improprio la batteria.
- Non accorciare i contatti esterni o gettare in acqua o sul fuoco la batteria.
- Sostituire la batteria solo con i ricambi HP previsti a questo scopo.

Le batterie e gli accumulatori non devono essere smaltiti insieme ai rifiuti domestici. Per procedere al riciclaggio o al corretto smaltimento, utilizzare il sistema di raccolta pubblico dei rifiuti o restituirli a HP, ai Partner Ufficiali HP o ai relativi rappresentanti.

Per ulteriori informazioni sulla sostituzione e sullo smaltimento delle batterie, contattare un Partner Ufficiale o un Centro di assistenza autorizzato.

Japanese battery notice

バッテリに関する注意

警告:本製品はバッテリを内蔵している場合があります。

- バッテリを取り外している場合は、充電しないでください。
- バッテリを水にさらしたり、60°C (140°F) 以上の温度にさらさないでください。
- バッテリを誤用、分解、破壊したり、穴をあけたりしないでください。
- 外部極を短絡させたり、火や水に投棄しないでください。
- バッテリを交換する際は、HP指定の製品と交換してください。

バッテリ、バッテリ パック、蓄電池は一般の家庭廃棄物と一緒に廃棄しないでください。リサイクルまたは適切に廃棄するため、公共の収集システム、HP、HPパートナー、またはHPパートナーの代理店にお送りください。

バッテリ交換および適切な廃棄方法についての情報は、HPのサポート窓口にお問い合わせください。

Spanish battery notice

Declaración sobre las baterías

ADVERTENCIA: Este dispositivo podría contener una batería.

- No intente recargar las baterías si las extrae.
- Evite el contacto de las baterías con agua y no las exponga a temperaturas superiores a los 60 °C (140 °F).
- No utilice incorrectamente, ni desmonte, aplaste o pinche las baterías.
- No cortocircuite los contactos externos ni la arroje al fuego o al agua.
- Sustituya las baterías sólo por el repuesto designado por HP.

Las baterías, los paquetes de baterías y los acumuladores no se deben eliminar junto con los desperdicios generales de la casa. Con el fin de tirarlos al contenedor de reciclaje adecuado, utilice los sistemas públicos de recogida o devuélvalas a HP, un distribuidor autorizado de HP o sus agentes.

Para obtener más información sobre la sustitución de la batería o su eliminación correcta, consulte con su distribuidor o servicio técnico autorizado.

Glossary

ACE	Access control entry.
ACL	Access control list.
ADS	Active Directory Service.
ALB	Advanced load balancing.
BMC	Baseboard Management Configuration.
CIFS	Common Internet File System. The protocol used in Windows environments for shared folders.
CLI	Command-line interface. An interface comprised of various commands which are used to control operating system responses.
CSR	Customer self repair.
DAS	Direct attach storage. A dedicated storage device that connects directly to one or more servers.
DNS	Domain name system.
FTP	File Transfer Protocol.
GSI	Global service indicator.
HA	High availability.
HBA	Host bus adapter.
HCA	Host channel adapter.
HDD	Hard disk drive.
IAD	HP 9000 Software Administrative Daemon.
iLO	Integrated Lights-Out.
IML	Initial microcode load.
IOPS	I/Os per second.
IPMI	Intelligent Platform Management Interface.
JBOD	Just a bunch of disks.
KVM	Keyboard, video, and mouse.
LUN	Logical unit number. A LUN results from mapping a logical unit number, port ID, and LDEV ID to a RAID group. The size of the LUN is determined by the emulation mode of the LDEV and the number of LDEVs associated with the LUN.
MTU	Maximum Transmission Unit.
NAS	Network attached storage.
NFS	Network file system. The protocol used in most UNIX environments to share folders or mounts.
NIC	Network interface card. A device that handles communication between a device and other devices on a network.
NTP	Network Time Protocol. A protocol that enables the storage system's time and date to be obtained from a network-attached server, keeping multiple hosts and storage devices synchronized.
OA	Onboard Administrator.
OFED	OpenFabrics Enterprise Distribution.
OSD	On-screen display.
OU	Active Directory Organizational Units.
RO	Read-only access.
RPC	Remote Procedure Call.
RW	Read-write access.
SAN	Storage area network. A network of storage devices available to one or more servers.
SAS	Serial Attached SCSI.

SELinux	Security-Enhanced Linux.
SFU	Microsoft Services for UNIX.
SID	Secondary controller identifier number.
SNMP	Simple Network Management Protocol.
TCP/IP	Transmission Control Protocol/Internet Protocol.
UDP	User Datagram Protocol.
UID	Unit identification.
VACM	SNMP View Access Control Model.
VC	HP Virtual Connect.
VIF	Virtual interface.
WINS	Windows Internet Naming Service.
WWN	World Wide Name. A unique identifier assigned to a Fibre Channel device.
WWNN	World wide node name. A globally unique 64-bit identifier assigned to each Fibre Channel node process.
WWPN	World wide port name. A unique 64-bit address used in a FC storage network to identify each device in a FC network.

Index

Symbols

/etc/sysconfig/i18n file, 28
9300 system
  components, 25
  configuration, 27
  features, 25
  management interfaces, 28
  shut down, 107
  software, 25
  start, 108
9320 system
  components, 25
  configuration, 27
  features, 25
  management interfaces, 28
  shutdown, 107
  software, 25
  start, 108

A

agile Fusion Manager, 53
AutoPass, 128

B

backups
  file systems, 77
  Fusion Manager configuration, 77
  NDMP applications, 77
battery replacement notices, 206

C

CLI, 32
clients
  access virtual interfaces, 50
cluster
  events, monitor, 94
  health checks, 95
  license key, 128
  license, view, 128
  log files, 98
  operating statistics, 98
  version numbers, view, 140
cluster interface
  change IP address, 125
  change network, 125
  defined, 122
components
  9300 diagrams, 180
  9320 diagrams, 183
contacting HP, 151
core dump, 68

D

Disposal of waste equipment, European Union, 202
document
  related information, 151
documentation
  HP website, 151
  providing feedback on, 153

E

email event notification, 70
events, cluster
  add SNMPv3 users and groups, 74
  configure email notification, 70
  configure SNMP agent, 72
  configure SNMP notification, 72
  configure SNMP trapsinks, 73
  define MIB views, 74
  delete SNMP configuration elements, 75
  enable or disable email notification, 71
  list email notification settings, 72
  list SNMP configuration, 75
  monitor, 94
  MSA array systems, 75
  remove, 95
  types, 70
  view , 94
exporting
  NFS, 154

F

failover
  automated, 50
  configure automated failover manually, 62
  crash capture, 68
  fail back a node, 64
  manual, 64
  NIC, 49
  server, 55
  troubleshooting, 141
Federal Communications Commission notice, 196
file serving node
  recover, 144
file serving nodes
  fail back, 64
  failover manually, 64
  health checks, 96
  maintain consistency with configuration database, 142
  migrate segments, 115
  monitor status, 93
  operational states, 94
  power management, 109
  prefer a user network interface, 124
  remove from cluster, 120
  rolling reboot, 109
  run health check, 142
  start or stop processes, 110
  statistics, 98
  troubleshooting, 140
  tune, 110

view process status, 110
file system
migrate segments, 115
firewall configuration, 34
firmware, upgrade, 129
Fusion Manager
agile, 53
back up configuration, 77
failover, 53

G

grounding
methods, 192
GUI
add users, 32
change password, 33
customize, 31
Details panel, 31
Navigator, 31
open, 29
view events, 94

H

hardware, power off, 108
hazardous conditions
symbols on equipment, 193
HBAs
display information, 66
monitor for high availability, 64
health check reports, 96
help
obtaining, 151
High Availability
agile Fusion Manager, 53
automated failover, turn on or off, 63
check configuration, 66
configure automated failover manually, 62
detailed configuration report, 67
fail back a node, 64
failover protection, 26
HBA monitor, 64
manual failover, 64
NIC HA, 54
power management for nodes, 109
power sources, 62
server HA, 54
summary configuration report, 67
troubleshooting, 141
host groups, 81
add domain rule, 82
add StoreAll client, 82
create host group tree, 82
delete, 83
prefer a user network interface, 124
view, 83
HP
technical support, 151
HP Insight Remote Support, 35
Phone Home, 37
troubleshooting, 46

I

Ibrix Collect, 134
add-on scripts, 137
configure, 136
ibrix_reten_adm -u command, 161
IP address
change for cluster interface, 125
change for StoreAll client, 125

L

labels, symbols on equipment, 193
laser compliance notices, 199
link state monitoring, 51
Linux StoreAll clients, upgrade, 18, 158
loading rack, warning, 193
localization, 28
log files, 98
collect for HP Support, 134

M

manpages, 33
monitoring
chassis and components, 88
cluster events, 94
cluster health, 95
file serving nodes, 93
node statistics, 98
servers , 84, 88
storage and components, 92

N

NDMP backups, 77
cancel sessions, 79
configure NDMP parameters, 78
rescan for new devices, 80
start or stop NDMP Server, 79
view events, 80
view sessions, 79
view tape and media changer devices, 80
network interfaces
add routing table entries, 126
bonded and virtual interfaces, 122
defined, 122
delete, 126
delete routing table entries, 126
guidelines, 48
viewing, 126
Network Storage System
configuration, 27
management interfaces, 28
NFS
exporting, 154
NIC failover, 49
no_subtree_check, 154
NTP servers, 35

P

passwords, change
  GUI password, 33
Phone Home, 37
ports, open, 34
power sources, server, 62
pre-6.3 Express Query, upgrade, 19

Q

QuickRestoreDVD, 144

R

rack stability
  warning, 152
recycling notices, 202
regulatory compliance
  Canadian notice, 197
  European Union notice, 197
  identification numbers, 196
  Japanese notices, 198
  Korean notices, 198
  laser, 199
  recycling notices, 202
  Taiwanese notices, 199
related documentation, 151
rolling reboot, 109
routing table entries
  add, 126
  delete, 126

S

segments
  evacuate from cluster, 118
  migrate, 115
servers
  configure standby, 49
  crash capture, 68
  failover, 55
  tune, 110
SNMP event notification, 72
SNMP MIB, 74
spare parts
  obtaining information, 152
Statistics tool, 100
  enable collection and synchronization, 100
  failover, 104
  Historical Reports GUI, 101
  install, 100
  log files, 106
  maintain configuration, 104
  processes, 105
  reports, 102
  space requirements, 103
  troubleshooting, 105
  uninstall, 106
  upgrade, 101
Storage
  software , 25
storage, remove from cluster, 118
StoreAll clients
  add to host group, 82
  change IP address, 125
  identify a user network interface, 123
  monitor status, 93
  prefer a user network interface, 124
  start or stop processes, 110
  troubleshooting, 140
  tune, 110
  tune locally, 114
  user interface, 33
  view process status, 110
StoreAll software
  shut down, 107
  start, 108
  upgrade, 10, 154
StoreAll software 5.5 upgrade, 167
StoreAll software 5.6 upgrade, 163
Subscriber's Choice, HP, 152
`subtree_check`, 154
symbols
  on equipment, 193
system recovery, 144
system startup, 108

T

technical support
  HP, 151
  service locator website, 152

U

upgrade60.sh utility, 159
upgrades
  6.0 file systems, 161
  firmware, 129
  Linux StoreAll clients, 18, 158
  pre-6.0 file systems, 159, 161
  pre-6.3 Express Query, 19
  StoreAll 5.5 software, 167
  StoreAll software, 10, 154
  StoreAll software 5.6 release, 163
  Windows StoreAll clients, 19, 159
user network interface
  add, 122
  configuration rules, 126
  defined, 122
  identify for StoreAll clients, 123
  modify, 123
  prefer, 123
  unprefer, 125

V

virtual interfaces, 48
  bonded, create, 49
  client access, 50
  configure standby servers, 49
  guidelines, 48

W

warning
rack stability, 152
warnings
loading rack, 193
websites
HP, 152
HP Subscriber's Choice for Business, 152
product manuals, 151
spare parts, 152
Windows StoreAll clients, upgrade, 19, 159